Innovative Technology for Computer Professionals

Computer

September 2002

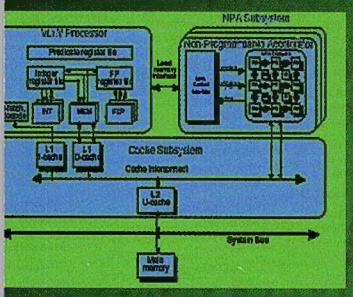

embedded SYSTEMS

Distributed Operating Systems

Web Services and Context Horizons

Jobs, Trades, Skills, and the Profession

IEEE

IEEE COMPUTER SOCIETY

http://computer.org

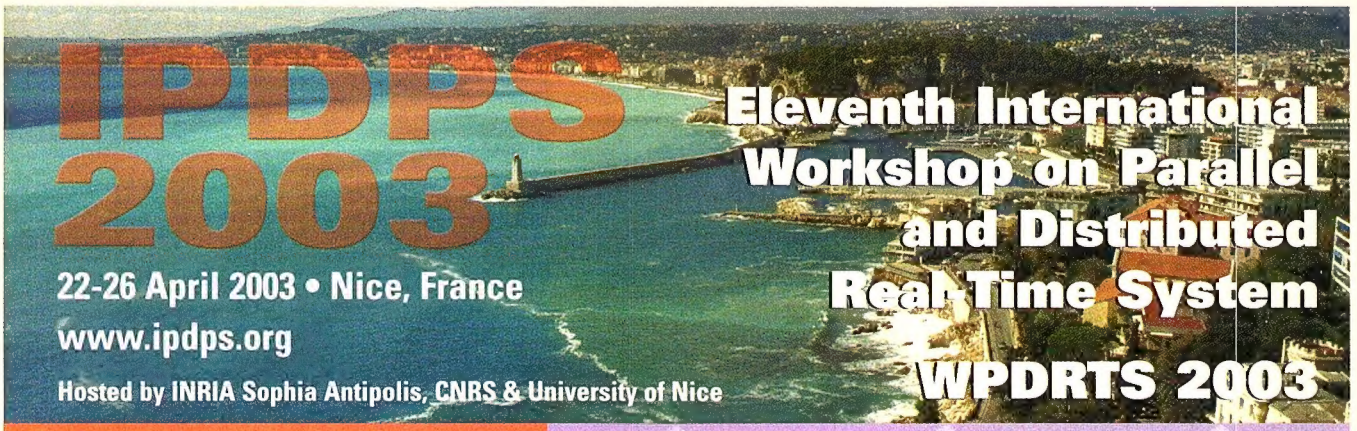

IPDPS 2003

Eleventh International Workshop on Parallel and Distributed Real-Time System

WPDRTS 2003

22-26 April 2003 • Nice, France

www.ipdps.org

Hosted by INRIA Sophia Antipolis, CNRS & University of Nice

IPDPS Workshops provide attendees an opportunity to spend a full day exploring special topics. Workshops planned for IPDPS 2003 are listed below. Each workshop has its own requirements and schedule for submissions, and all are linked from the IPDPS Web site at www.ipdps.org.

- Heterogeneous Computing
- High-Level Parallel Programming Models & Supportive Environments
- Java for Parallel and Distributed Computing
- Parallel and Distributed Computing in Image Processing, Video Processing, and Multimedia
- Advances in Parallel and Distributed Computational Models
- Reconfigurable Architectures
- Communication Architecture for Clusters
- NSF Next Generation Systems Program
- High Performance Computational Biology
- Wireless, Mobile, and Ad Hoc Networks
- Fault-Tolerant Parallel and Distributed Systems
- Nature Inspired Distributed Computing
- Formal Methods for Parallel Programming
- Internet Computing and E-commerce
- Parallel and Distributed Scientific and Engineering Computing with Applications
- Massively Parallel Processing
- Performance Modeling, Evaluation, and Optimization of Parallel and Distributed Systems
- Massively Parallel Computer Aided Surgery
- Parallel and Distributed Systems: Testing and Debugging

IPDPS 2003
IPPS & SPDP

International Parallel & Distributed Processing Symposium

Sponsored by the
IEEE Computer Society Technical Committee on Parallel Processing

In cooperation with ACM SIGARCH, IEEE Computer Society Technical Committee on Computer Architecture (TCCA), and IEEE Computer Society Technical Committee on Distributed Processing (TCDP)

WPDRTS 2003 will be held in Nice, France on April 22nd & 23rd in conjunction with IPDPS 2003. This workshop is a forum for the presentation and discussion of approaches, research findings, and experiences in the domain of large-scale parallel and distributed real-time systems. Of interest are both the research and development of relevant technologies as well as the applications built using such technologies.

KEYNOTE SPEAKERS

Inter-Vehicle Real-time Communication and Networking
Dr. Ralf G. Herrtwich, Director Telematics Research, DaimlerChrysler
Real-time Flight Software Systems of Today and Tomorrow
Ken Rehm, NASA Goddard Space Flight Center

CHALLENGE PROBLEM

Michael Masters, Chief Scientist of the High Performance Distributed Computing Program of the US Navy, has provided a challenge problem on the topic: *System Certification for Real-time Systems that Employ Dynamic Resource Management.* Proposed solutions to the challenge problem are solicited. A special session will be devoted to discussing the problem and the proposed solutions. (For details, see http://wpdrts.cs.ohiou.edu/challenge.html.) An award will be given for the best solution, and the best solution will be printed in *The Journal of Systems and Software.*

CALL FOR PAPERS

Topics of interest include but are not limited to: resource management; operating systems and middleware; programming environments; algorithms and applications; architectures; specification, modeling, and analysis; & networking and communications. Papers are not to exceed 15 pages (11 pt. font). Please number each page. Include an abstract, 5-10 keywords, and the technical areas most relevant to your paper. The workshop proceedings will be published by IEEE CS Press, and selected papers will be published in *The Journal of Systems and Software.* Submission procedures are available via the WPDRTS 2003 Web site at **http://wpdrts.cs.ohiou.edu**. **Papers are due November 1, 2002**. Acceptance notification will be sent by December 20, 2002 and final manuscripts are due January 20, 2003. For further information, contact info@wpdrts.cs.ohiou.edu.

WPDRTS 2003 ORGANIZATION

Steering Committee Co-Chairs
David Andrews, *University of Kansas*
Lonnie R.Welch, *Ohio University*

General Co-Chairs
Barbara Pfarr,
NASA Goddard Space Flight Center
Armin Zimmerman,
Technische Universität, Berlin

Program Co-Chairs
Klaus Ecker,
Institut für Informatik, Germany
Eui-Nam Huh,
Sahmyook University, Korea
G. Manimaran, *Iowa State University*

Steering Committee
Scott Brandt,
University of California at Santa Cruz
Alan Burns, *The University of York, UK*
Guenter Hommel,
Technische Universität, Berlin, Germany
Doug Locke, *Timesys Corporation*
Viktor Prasanna,
University of South California
Behrooz Shirazi,
University of Texas at Arlington
Kenji Toda,
Electrotechnical Laboratories, Japan
Wei Zhao, *Texas A&M University*

The Latest in Computing Technology From

WILEY IEEE

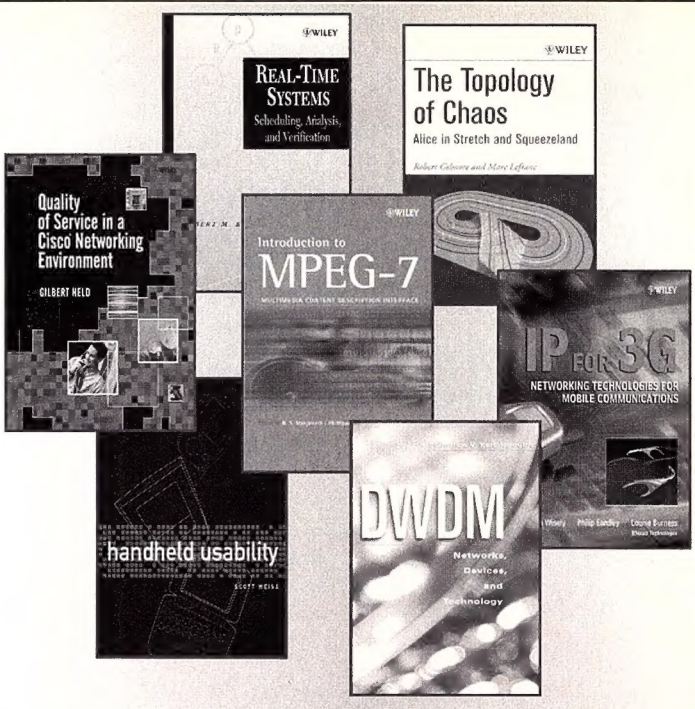

From John Wiley & Sons

Handheld Usability
Scott Weiss

Handheld Usability is a practical, hands-on guide to designing interfaces for handheld, electronic computing and communication devices, including e-mail pagers, Personal Digital Assistants (PDAs) and mobile telephone handsets.

0-470-84446-9 2002 292pp. $55.00

Introduction to MPEG-7
Multimedia Content Description Interface
B.S. Manjunath, Philippe Salembier, Thomas Sikora

Takes a systematic approach to provides a unique overview of the principles and concepts behind audio-visual indexing, metadata description, information retrieval and browsing.

0-471-48678-7 2002 396pp $90.00

Quality of Service in a Cisco® Networking Environment
Gilbert Held

Guides the reader through the theory behind a specific QoS technique and is followed with a series of command examples tailored to a specific QoS technique. Readers learn how a QoS technique operates and how use it in a Cisco environment.

0-47084425-6 2002 200pp. $55.00

IP for 3G Networking Technologies for Mobile Communications
Dave Wisely, Philip Eardley, Louise Burness

IP for 3G gives a comprehensive overview of 3G networking functionality and examines how IP protocols can be developed to provide some of the basic building blocks of a mobile system (mobility, QoS and call control).

0-471-48697-3 2002 304pp. $55.00

Real-Time Systems Scheduling, Analysis, and Verification
Albert M. K. Cheng

Provides a substantial, up-to-date overview of the verification and validation process. It describes several approaches to the problem, and examines the advantages and disadvantages of each approach. Ample illustrations and examples of a variety of industrial and toy applications, it is a key resource for every professional who works with real-time systems.

0-471-18406-3 2002 552pp. $95.00

The Topology of Chaos
Alice in Stretch and Squeezeland
Robert Gilmore, Marc Lefranc

This volume responds to the fundamental challenge of chaotic systems by introducing a new analysis method—Topological Analysis. It can be used to extract from chaotic data the topological signatures that determine the stretching and squeezing mechanisms which act on flows in phase space and are responsible for generating chaotic data.

0-471-40816-6 2002 495pp. $84.95

From Wiley-IEEE Press

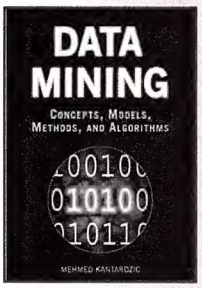

Data Mining Concepts, Models, Methods and Algorithms
Mehmed Kantardzic

This book offers guidance to selecting among the numerous analysis programs. It discusses data mining principles and representative state-of-the-art methods and algorithms originating from different disciplines such as statistics, machine learning, neural networks, fuzzy logic, and evolutionary computation. Detailed algorithms are provided with necessary explanations and illustrative examples.

0-471-22852-4 Nov. 2002 360pp. $59.45

Power System Control and Stability, 2nd Edition
Paul M. Anderson

Fully updated and expanded to include the latest developments in the field, it describes the mechanical system that drives the electric generators, and the dynamic reaction between the prime mover and generator systems.

0-471-23862-7 Nov. 2002 672pp. $99.95

DWDM Networks, Devices, and Technology
Stamatios V. Kartalopoulos

Provides a comprehensive treatment of DWDM, technology, systems, and networks, as well as engineering design. It explains how DWDM works, how to use DWDM in system design, how optical network architecture can benefit from DWDM.

0-471-26905-0 Oct. 2002 520pp. $74.95

TO ORDER:
MAIL: Customer Care-Wiley, 10475 Crosspoint Blvd., Indianapolis, IN 46256
CALL: 1-877-762-2974 **EMAIL:** custserv@wiley.com **FAX:** 1-800-597-3299
For orders outside of North America:
Call: +44 (0) 1243-843-2294 **Fax:** +44 (0) 1243 -843-296.
Email: cs-books@wiley.co.uk
Prices subject to change without notice and may be higher outside of U.S.
Appropriate local sales tax and shipping and handling will be added to all orders.

Innovative Technology for Computer Professionals
Computer
September 2002, Volume 35, Number 9

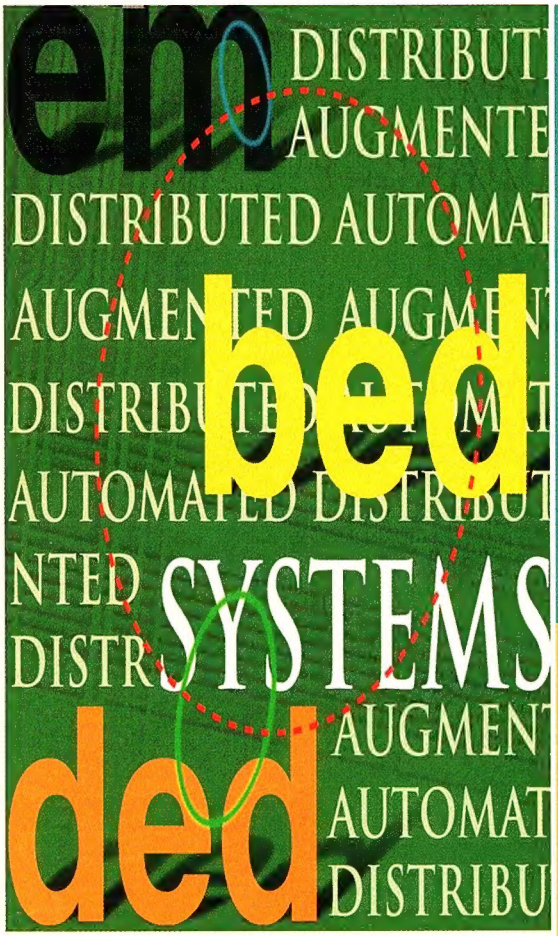

Cover design and artwork by Dirk Hagner

ABOUT THIS ISSUE

Embedded computers are everywhere—in video games, TV sets, printers, cell phones, and cars. The advent of system-level integration foreshadows an era of yet more growth in the number and variety of innovative smart products. In this issue, we look at the PICO synthesis system for automatically designing custom computers and a sample next-generation application—a smart camera—that represents a quantum leap in sophistication and pushes the embedded design space in multiple directions. Other articles look at an augmented reality system that gives paleontologists new insights into their research, branching strategies for software configuration management, and the design of mobile agent communication protocols.

Computer
Innovative Technology for Computer Professionals

Circulation: *Computer* (ISSN 0018-9162) is published monthly by the IEEE Computer Society. **IEEE Headquarters**, Three Park Avenue, 17th Floor, New York, NY 10016-5997; **IEEE Computer Society Publications Office**, 10662 Los Vaqueros Circle, PO Box 3014, Los Alamitos, CA 90720-1314; voice +1 714 821 8380; fax +1 714 821 4010; **IEEE Computer Society Headquarters**, 1730 Massachusetts Ave. NW, Washington, DC 20036-1903. IEEE Computer Society membership includes $14 for subscription of *Computer* magazine ($14 for students). Nonmember subscription rate available upon request. Single-copy prices: members $10.00; nonmembers $20.00. This magazine is also available in microfiche form.

Postmaster: Send undelivered copies and address changes to *Computer*, IEEE Service Center, 445 Hoes Lane, Piscataway, NJ 08855. Periodicals Postage Paid at New York, New York, and at additional mailing offices. Canadian GST #125634188. Canada Post Publications Mail (Canadian Distribution) Agreement Number 0487910. Printed in USA.

Computer

IEEE Computer Society: http://computer.org
Computer: http://computer.org/computer
computer@computer.org
IEEE Computer Society Publications Office: +1 714 821 8380

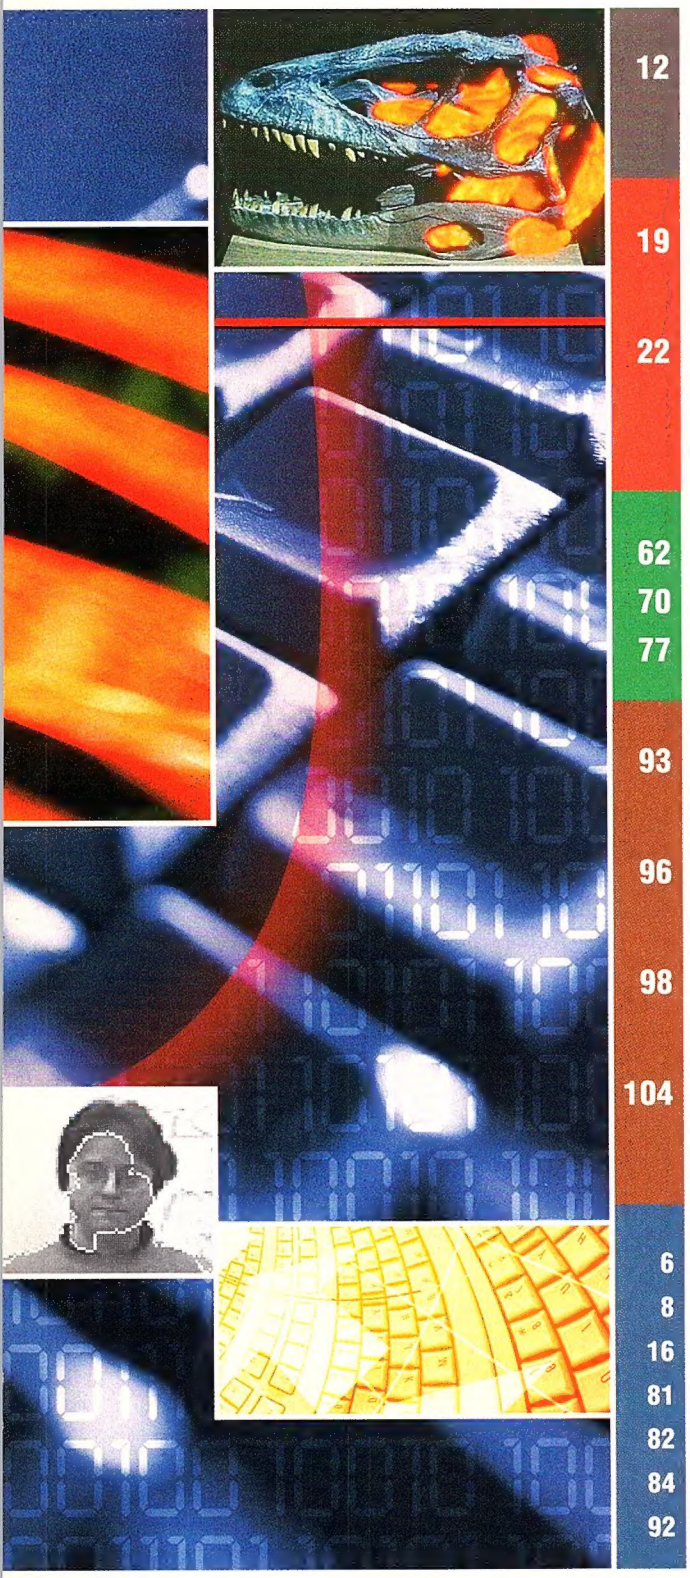

Membership Magazine of the IEEE COMPUTER SOCIETY

IEEE

ARTICLE SUMMARIES

Merging Fossil Specimens with Computer-Generated Information

pp. 25-30

Oliver Bimber, Stephen M. Gatesy, Lawrence M. Witmer, Ramesh Raskar, and L. Miguel Encarnação

While dinosaur fossils reveal much, they keep us guessing about the original organisms' color, sound, and behavior. For several years, modern paleontologists have used 3D computer graphics to help reconstruct these pieces of the past.

Augmented reality leverages this technology to provide an interface that enhances the real world with synthetic supplements. Paleontologists can use AR to present virtual data directly within a real, 3D environment rather than on a flat monitor. This technology's immersiveness can give paleontologists new insights into their research and communicate the results to museum visitors in an exciting and effective way.

The Importance of Branching Models in SCM

pp. 31-38

Chuck Walrad and Darrel Strom

To improve software quality, you must first understand your software. If you do not understand your code base, your odds of updating it without breaking something are poor. Often, a fundamental misunderstanding of software configuration management (SCM) as it applies to real-world application development is at fault.

Branching is integral to version management, software build correctness, and release management. Good decisions about when and why to branch can make it much easier for developers and release engineers to coordinate software product changes. The right branching strategy makes it easier to deliver the right code, re-create past releases, and—if necessary—roll back to a previous release.

Adopting the right SCM branching model facilitates rapid development, increases overall product quality and process efficiency, reduces the incidence of software failures, and improves organizational performance.

PICO: Automatically Designing Custom Computers

pp. 39-47

Vinod Kathail, Shail Aditya, Robert Schreiber, B. Ramakrishna Rau, Darren C. Cronquist, and Mukund Sivaraman

The PICO (program in, chip out) project is a long-range HP Labs research effort that aims to automate the design of optimized, application-specific computing systems—thus enabling the rapid and cost-effective design of custom chips when no adequately specialized, off-the-shelf design is available.

PICO research takes a systematic approach to the hierarchical design of complex systems and advances technologies for automatically designing custom nonprogrammable accelerators and VLIW processors. While skeptics often assume that automated design must emulate human designers who invent new solutions to problems, PICO's approach is to automatically pick the most suitable designs from a well-engineered space of designs. Such automation of embedded computer design promises an era of yet more growth in the number and variety of innovative smart products by lowering the barriers of design time, designer availability, and design cost.

Smart Cameras as Embedded Systems

pp. 48-53

Wayne Wolf, Burak Ozer, and Tiehan Lv

Recent technological advances are enabling a new generation of *smart cameras* that represent a quantum leap in sophistication. While today's digital cameras capture images, smart cameras capture high-level descriptions of the scene and analyze what they see. These devices could support a wide variety of applications including human and animal detection, surveillance, motion analysis, and facial identification.

Video processing has an insatiable demand for real-time performance. Smart cameras leverage very large-scale integration to meet this need in a low-cost, low-power system with substantial memory. Moving well beyond pixel processing and compression, these VLSI systems run a wide range of algorithms to extract meaning from streaming video.

Recently, Princeton University researchers developed a first-generation smart camera system that can detect people and analyze their movement in real time. Because they push the design space in so many dimensions, these smart cameras are a leading-edge application for embedded system research.

Mailbox-Based Scheme for Designing Mobile Agent Communication Protocols

pp. 54-60

Jiannong Cao, Xinyu Feng, Jian Lu, and Sajal K. Das

In various situations, mobile agents at different hosts must cooperate with one another by sharing information and making decisions collectively. To ensure effective interagent communication, communication protocols must track target agent locations and deliver messages reliably.

Researchers have proposed a wide range of schemes for agent tracking and reliable message delivery. However, each scheme has its own assumptions, design goals, and methodology. As a result, no uniform or structured methods exist for characterizing current protocols, making it difficult to evaluate their relative effectiveness and performance.

The authors propose a mailbox-based scheme for designing mobile agent communication protocols. This scheme assigns each agent a mailbox to buffer messages, but decouples the agent and mailbox to let them reside at different hosts and migrate separately.

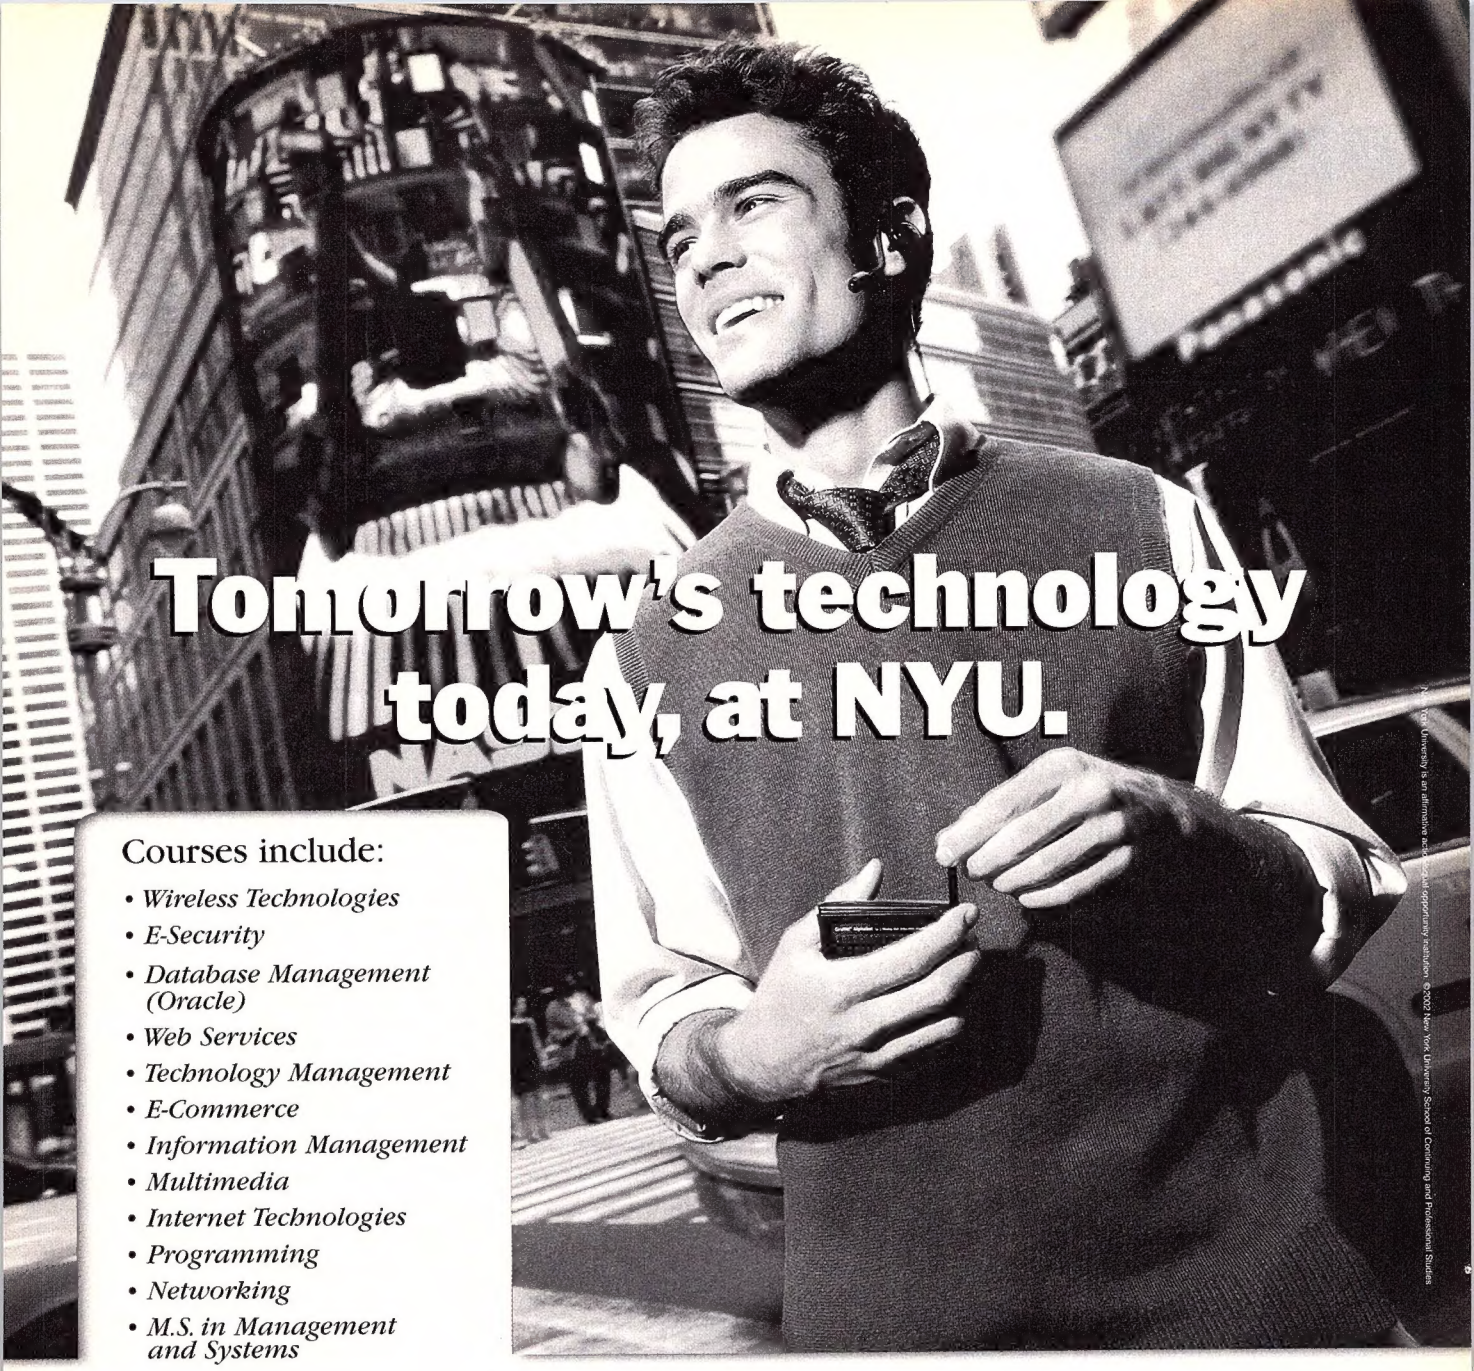

Courses include:

- *Wireless Technologies*
- *E-Security*
- *Database Management (Oracle)*
- *Web Services*
- *Technology Management*
- *E-Commerce*
- *Information Management*
- *Multimedia*
- *Internet Technologies*
- *Programming*
- *Networking*
- *M.S. in Management and Systems*

The only thing changing faster than today's world is today's technology. That's why NYU's School of Continuing and Professional Studies has developed programs designed to keep you at the forefront of this rapidly evolving industry. We offer a broad range of courses, certificates, and degrees, including a powerful Master's Degree, all taught by the industry specialists who are leading the changes.

From enterprise-wide solutions to e-security, from technology management to wireless, tap into the future, and grab ahold of yours. Find out more – in a nanosecond – before anything else changes! Call us or visit our website.

NEW YORK UNIVERSITY
A PRIVATE UNIVERSITY IN THE PUBLIC SERVICE
School of Continuing and Professional Studies

For more information	**Phone:** 1-800-FIND NYU, ext.192	**Website:** www.scps.nyu.edu/x192

Home Theater Strategies

To the Editor:

As Donald Norman points out in "Home Theater: Not Ready for Prime Time" (June 2002, pp. 100-102), installing a home theater is beyond most users' ability. Sorting through the array of products that are not compatible, either mechanically or aesthetically, requires significant knowledge of both theory and products.

I have both a BS and an MS in computer science, am proficient at programming, and am adept at mechanical things, yet setting up my home theater was still an exceedingly difficult task. However, after a great deal of study and effort, I finally acquired all the components that fit together and actually sound good, and I now have a home theater that I enjoy immensely.

I propose that manufacturers should make their components compatible with computer and Internet communications such as an Ethernet Network Interface Card and an RJ45 Ethernet jack, so that users could update their components from their TV screen or computer.

Users could accomplish simple programming, such as identifying inputs and outputs, on a TV screen. They could use a browser interface on their computer to perform more complex programming, such as pressing a single button to have the system go into a particular mode—as in switch TV on, place TV in DVD mode, turn on DVD, turn on sound system, tune sound system in to DVD.

This strategy would further the current trend among vendors to move from a hardware-based approach to a

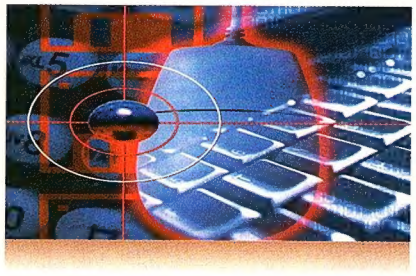

software-based alternative. Vendors would only need to make a relatively small change to provide a common interface to their components, instead of spending the time and effort required to agree upon industry standards for their various hardware and software interfaces.

Although it would not solve all of the problems, this solution would make it possible for a typical user to set up—but not necessarily install—a home theater with minimum effort.

Ray Kolbe
Santa Barbara, Calif.
rkolbe@computer.org

PROFESSIONAL FOLLIES

To the Editor:

Bob Colwell's "Engineering Follies" (June 2002, pp. 10-12) made me laugh and brought back memories of my own undergraduate days.

My MS is in computer engineering, but I like EE, and almost became one. As a mechanical engineering student, I had to take the usual three EE courses during my junior year. We used smaller motors in our lab, about one-half meter long and about the same diameter. For loads, we used banks of huge lightbulbs that were wired together with coat-hanger-sized wire and switched with old-fashioned knife switches.

One day when we were doing the DC motor lab, the instructor told us a story about a shipboard motor that had gone for a flight after the field opened, crashing through a steel bulkhead and killing someone on the other side. With this story fresh in our minds, we began switching on the load banks.

Unfortunately, the student running the switches moved a bit slowly, and the switch arced and started sparking. Shouts of "Shut it off!" sent all the students running for the door just as the instructor, who had left the room to get a drink of water, was trying to get back in. Later, when the chaos had settled down, one of the wires came loose, but by then shouting "Shut it off!" had a less exciting effect.

I tell this story to my students when I teach the same lab at Cal Poly. That I can still remember it proves Bob's point that "When the real world intervenes on well-planned experiments, the lessons you learn stick with you." Amen.

I add a similar story to each of my lab assignments, trying to match the subject or lesson to the particular experiment. The students love it, and it actually entices some to read the rest of the lab.

I feel strongly that societies such as the IEEE and ASME should collect such stories from their members. I fear that we may lose a great anecdotal history of our profession if we fail to do this.

Mark S. Hutchenreuther
San Luis Obispo, Calif.
mhutchen@falcon.csc.calpoly.edu

UNCOVERING INVASIVE SOFTWARE

To the Editor:

George Lawton's article on spyware ("Invasive Software: Who's Inside Your Computer?" July 2002, pp. 15-18) was particularly helpful to me. Thanks to this article and Lavasoft, I uncovered the hidden software that has saturated our company's IT system.

In addition to being the MIS director for my company, I am also the net-

work administrator and IT manager. To properly manage the many responsibilities I am tasked with, I must digest as much information as possible relating to current trends.

Although some of what I read in *Computer* is over my head, articles such as this one that address fundamental technological issues help to remind me of the value of my IEEE Computer Society membership.

William D. Jackson
Atlanta, Ga.
wdjacks@emoryfcu.com

CERTIFICATION DEBATE

To the Editor:

Although I appreciate the attempt to present a substantive debate about software certification, I feel that Adam Kolawa ("Certification Will Do More Harm than Good," June 2002, pp. 34-35) lacked objectivity in his criticism of the IEEE Computer Society's proposed Certified Software Development Professional (CSDP) program.

Kolawa claims that certification will stifle innovation by "making developers more rigid." A mature software development process and a software development certification program actually help to nurture innovation by allowing developers to focus on the areas that require creativity. An organization that hires "natural" programmers who reject any process or discipline is bound to suffer the consequences as these programmers take their tacit knowledge with them when they decide to leave.

The CSDP does not intend to test developers on fundamentals that they should have learned in an accredited CS program. Rather, the program exposes developers to practical software development practices. The CSDP documentation does a nice job of describing the differences between accreditation, certification, and skill development and explaining their interconnection.

Kolawa struggles with the fact that computer science is actually considered a science. Simply stated, science is the interplay of logic with observation, a definition that clearly applies to computer science. I believe that computer science research is more complex than research in the pure sciences because it addresses both quantitative and qualitative components. In addition, it is systemic as opposed to systematic as in the social sciences—which also have certification programs.

Barry M. Baker
Fishkill, N.Y.
barrymbaker@ieee.org

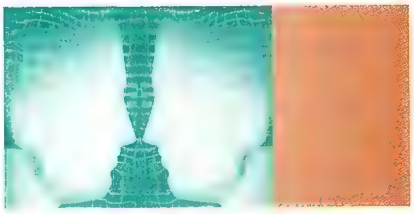

To the Editor:

After reading "Benefits of Certification" (June 2002, pp. 31-33), I concluded that Leonard Tripp must work in a research department at the Boeing Co., not in a typical corporate IT/IS department.

Most corporate IT departments and software development firms have knowledgeable, fully qualified software engineers on staff. However, department managers and senior software engineers often face unrealistic deadlines for fixing outstanding issues or releasing the next revision of the software they are presently working on.

All too often, the pressure to meet these deadlines forces the software engineers to take shortcuts and write "spaghetti code." This pressure causes a ripple effect from the beginning stages of development to quality testing before the code is released.

Concerning the statement that certification may survive budget cuts in difficult economic times, I have not found this to be true. For the most part, surviving budget cuts depends more on toeing the company line and the old saying, "It's not what you know, but who you know."

Tripp mentions a "lack of awareness of accepted best practices," but this is subjective because what is good for one person or project is not necessarily good for another.

However, if this statement means that today's software engineers are not learning the appropriate principles when they pursue their bachelor's and master's degrees from accredited universities, these programs should address this issue. After all, this is where we "create" tomorrow's software engineers.

Todd Kolb
Mahwah, N.J.
thekolbs@optonline.com

To the Editor:

In "Benefits of Certification," Leonard Tripp states that "users expect software developers to meet certain standards of competence and ethics." In my experience, when someone asks about ethics, it's often because they are suggesting something unethical.

In this context, the definition of

"user" makes a significant difference. Is Tripp referring to the people who buy the software or to our employers?

Currently, there are obvious differences between working in a large organization that has adopted formal methods and has established predefined and structured procedures for verifying the work and working at a small company where the programming is more like hacking. If software development is to fully become an engineering profession, we must establish methods that everyone can use, without regard to differences in their working environment.

Certification can have different meanings depending on the company that is doing the certifying. Worse, if the certification process deviates too much from an academic course of study, it is impossible to be certain that a software developer has received certification because of his knowledge of the discipline or because he did well at taking a test. For certification to have a positive impact, we must make certain that it does not appear to be something an aspiring software developer can purchase at a shop on the corner.

Franco Favento
Trieste, Italy
f.favento@ieee.org

To the Editor:

While I agree with most of Leonard Tripp's comments on certification, I disagree with his statement that "Certification provides a tool for measuring knowledge and skills so that the public can understand an individual's knowledge and competence in relation to others having a similar background." This is an apparent contradiction with his comment that "Certification is not a guarantee of competency, but rather an assessment indicating that an individual understands a certain level of professional practice."

It would be more accurate to say that certification provides a tool for measuring a defined subset of knowledge and skills. In general, this subset probably would not include the creative skills needed for software design. A certification exam is more likely to measure "auditing" skills—the understanding needed to review and critique existing software or designs. Good auditing skills are important for creative work, but they do not guarantee programmer productivity.

I object to almost every statement that Adam Kolawa makes in "Certification Will Do More Harm than Good." Kolawa contends that certification would discount the importance of a "brilliant" 21-year-old programmer who may not have a degree or certification. However, I cannot imagine why a programmer who can write a system composed of say C++, XML, XSL, and .NET could not pass an exam in each of these subjects with little additional study.

Certification for software developers targets degreed individuals who have been in the field for 10 to 20 years or more and have little time to attend formal classes. These experienced developers have already demonstrated their creative competence on the job. They are primarily concerned with flexible self-study and a standardized testing system that confirms their competence and gives them feedback on how thoroughly they understand a particular area.

I doubt that potential employers would use certification test results to the exclusion of other indicators of professional success, any more than graduate schools would use GRE scores as the sole criterion for admissions. At the same time, standardized tests are a couple of orders of magnitude better than the technical questions I've encountered in interviews.

Bert Craytor
Pacifica, Calif.
craytor@montaramedia.com

To the Editor:

I agree with Adam Kolawa with regard to the limitations of software certification.

My experience teaching an upper-level course in computer information systems for more than 12 years clearly demonstrates that mental organization is the determining factor in a student's success.

I emphasize to my students that it is their ability to think effectively that matters. However, I often find that they are content to believe that things just happen—they aren't interested in knowing why. They have a "just write it and be done with it" mentality.

Unfortunately, certifications such as MSCE, Novell, and Cisco foster this kind of attitude because they emphasize memorization and learning by rote. But computer science is an art, and success in this field depends on attitude, insight, organized thought processes, logic, and hands-on skills.

Louis H.R. Muller
Centereach, N.Y.
lhrm@juno.com

We welcome your letters. Send them to computer@computer.org. Letters are subject to editing for style, clarity, and length.

PROGRAM AND
REGISTRATION MATERIAL
AVAILABLE NOW ONLY ON THE WEB

BALTIMORE SC2002 NOV 16-22

SPONSORS
IEEE Computer Society
ACM Sigarch

BALTIMORE CONVENTION CENTER
BALTIMORE, MARYLAND, USA

www.sc-conference.org/sc2002

SC2002 ADVANCE PROGRAM

FROM TERABYTES TO INSIGHTS

The conference that brings together scientists, engineers, educators, visualization artists, programmers, and managers from industry, academia, and government promises to be better than ever this year. In Baltimore, technical program participants and exhibitors from the U.S., Canada, Mexico, South America, Europe, and Asia will gather to showcase breakthroughs in information architectures, scalable computing, data analysis applications, visualization, and networking. SC2002 will also provide a vision of the changes that will impact our world in the next decade.

Thanks to high performance computing and networking, terabytes of data are being transformed into knowledge—about the origins of our universe, our genetic makeup, the condition of our ecosystems, and the challenges of security in a dangerous world.

GET SMART

For the first time in conference history, SC2002 technical program participants can take part in IntelliBadge™, a demonstration of "smart" technology. IntelliBadge participants will track events that fit their needs and interests, locate people of similar interests, view statistics on the conference, and win valuable prizes.

- **KEYNOTE ADDRESS:** Rita Colwell (invited), Director, National Science Foundation

TECHNICAL PROGRAM

- **PAPERS:** 67 exceptional papers presented out of 231 submitted
- **TUTORIALS:** 31 full and half-day tutorials chosen from a record 87 submissions
- **MASTERWORKS:** Sessions over three days in four general areas: infrastructure, life sciences, computer aided engineering, and HPC in commerce
- **PANELS:** Lively discussions by the experts on the hottest topics, including two panels on the use of information technology for homeland security
- **STATE OF THE FIELD TALKS:** Visionaries and well-known leaders offer their insights on the issues
- **AWARDS:** Presentation of the Gordon Bell prizes, Sidney Fernbach Memorial Award, Seymour Cray Computer Engineering Award, and the High Performance Bandwidth Challenge, sponsored by Qwest

EXHIBITION EVENTS

- **GALA GRAND OPENING:** A preview of the exhibits on Monday, 7-9 p.m.
- **INDUSTRY AND RESEARCH EXHIBITS:** More variety and more excitement from over 140 exhibitors
- **RESEARCH POSTERS:** Highlights of timely work by top researchers
- **EXHIBITOR FORUM:** Leaders in industry describe recent breakthroughs and new technical plans

EDUCATION PROGRAM

Hands-on sessions for teacher and faculty teams to help them integrate the latest computational tools and methods into the classroom

INTELLIBADGE DEMONSTRATION

A chance to put "smart" technology to use to enhance the conference experience

MINORITY SERVING INSTITUTIONS PROGRAM

Faculty and staff from Minority Serving Institutions broaden and enrich the SC community

SCinet

The conference's high-performance, production-quality network backbone that uses the most advanced technology to make available more than 55 Gbits/second of wide area network connectivity to the national networks.

Registration now open at
www.sc-conference.org/sc2002

Advance registration closes
October 11, 2002

Deadline for housing at conference rate
October 11, 2002

Employee Performance Reviews

Bob Colwell

Every year, many companies engage in employee performance reviews, a months-long process that scares their employees, costs huge amounts of management time, and purposely sunders the teamwork required for product development. From this process come promotions, raises, management changes, and, ultimately, much of the corporate culture itself. Performance reviews also generate demotions, disciplinary actions, and dismissals.

A substantial measure of overall job satisfaction derives from this official evaluation. Mix in all the normal human frailties about who likes whom better, misremembered communication mixups in meetings, incomplete or erroneous understandings about the evaluation process, and pretty soon the trepidation approaches mandatory-Valium levels.

Why do companies pursue this painful process every year? Because they believe in the idea of a technical meritocracy. Companies want to reward their top producers, and they need to identify those who may not be carrying their share of the load. Employees hope to distinguish themselves by their efforts, and they want those efforts to be noticed and rewarded.

People who contribute more to the company's efforts should be rewarded commensurately. Recognizing its best producers and rewarding them with promotions and bonuses encourages the kind of behavior the company wants to foster.

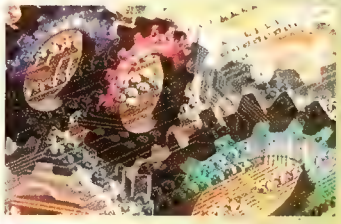

An R&R session groups the technical ranks according to their individual accomplishments.

Without performance reviews, deadwood tends to accumulate—people who are not enjoying their jobs, are not producing at a rate comparable to their peers, and probably are holding the entire team back from its goals. The question is, how can a company evaluate a population of engineers in such a way that they accept the process as fair and equitable, encouraging maximum overall output?

Since it is the example I know best and because I think its system works well overall, I will use Intel as my case study for how one company answers this question.

INTEL'S MERITOCRACY

Intel divides its employees into grades. A recent college graduate with a BSEE might start out as a grade 3. As the years go by and the employee gains in knowledge, experience, and responsibility, she typically will be promoted through the various grades, perhaps up to principal engineer (grade 10) or Intel Fellow (grade 12).

Promotions are determined at management meetings called "ranking and rating" sessions. The basic idea of an R&R is to divide up the technical ranks into groups of about 20 engineers and rank order the names according to their individual accomplishments over the past 12 months. The managers then use this rank order to drive discussions about pay raises, stock option allocations, promotions or demotions, and "messages"—the advice that will appear in the final written performance review each employee receives.

At these R&R sessions, an immediate supervisor represents each employee who is being ranked and rated. A department manager or the general manager (the "rank manager") leads the session. Typically, a human resources representative sits in to make sure the session follows Intel policy and to serve as an independent witness that all employees are treated fairly.

After the R&R, the supervisors have a few weeks in which to write the formal reviews, which include the accomplishments, strengths, and areas for improvement, along with whatever advice the supervisor feels will best guide the employee toward higher productivity and more job satisfaction. The supervisor, the supervisor's boss, and the employee sign these reviews. A separate letter details any salary, bonus, stock option, or pay grade changes.

NO SMOKY BACK ROOMS

Well, that's how things are supposed to work. It's a pretty logical process, and it usually arrives at a satisfactory result. However, at its core, this process does require humans to subjectively evaluate who did what, who gets credit for what,

how important various accomplishments are relative to one another, and so on. This is especially tricky when people have collaborated intensively on highly productive teamwork, which the company generally encourages.

The logistics of running the R&R session are extremely important, and there is wide variation across Intel on the particulars. I know, because our microprocessor design group in Oregon tried several different schemes before settling on what I think was the best format.

Scheme 1: Filibuster

When we used this method, supervisors arrived at the R&R armed with a set of PowerPoint slides—there were no online presentations in 1992. Every employee being ranked and rated was represented by a set of slides that the other supervisors had not seen prior to the R&R.

The session started with the rank manager asking for nominations for the #1 spot on the list. The supervisors proposed three or four names, and their slide shows began. With no time limit set in advance, the supervisors were free to extol the virtues of their employees at great length. This wasn't so bad at the beginning of the R&R, but it was excruciating by the end. No one could remember anything about the candidate from two slide shows ago, yet each candidate had to be ranked against the person currently being talked up.

This method was particularly worrisome to technical supervisors when the ranking group included nontechnical or quasi-technical people. Technical supervisors are not always great presenters, but marketing people often are. When a marketing supervisor is describing an employee, the technical supervisors are squirming in their seats, concerned that their engineers may be underappreciated solely because their representative lacks panache. Marketing supervisors worry that other supervisors may not appreciate the nuances of marketing and how difficult it is.

I participated in some R&R sessions with no fixed time limits per presentation and mixed ranks (design engineers mixed with marketing folks) that lasted for six hours or more, with no one feeling particularly good about the outcome when the session finally ended. As Intel would say: Improvement required.

A performance review structure simply doesn't work well when run as laissez-faire capitalism.

Scheme 2: Time limits

As in Scheme 1, supervisors arrived at the R&R with slide presentations representing their employees' numerous contributions to the company. But they had only three minutes for each presentation, which had the salutorious effect of limiting the overall length of the session, as well as helping normalize the marketing-engineering presenter gap.

That helped, but problems still remained. You can only hear and retain so much at any meeting. However, people are willing to do prep work ahead of the R&R if they can figure out how to arrange it, which led to Scheme 3.

Scheme 3: Preordering

A week before the R&R, every supervisor was required to turn in one slide per employee in a stylized format, the "333": name, job description, three biggest accomplishments, three strengths, and three areas for improvement. These slides were distributed to all supervisors, who used them to tentatively rank the employees based solely on what they read. The rank manager then collected these rankings and fed them into a spreadsheet to create an overall preordering.

The R&R started with this preordering. Many people were surprised to find how close to the final ordering this aggregate preordering was. Instead of wasting time arguing about the obvious, this method led directly to productive discussions.

Scheme 3 was a big win, but the process still required improvement.

Scheme 4: Fixed-font preordering

If you've written technical papers for conferences, you'll recognize this next gambit. Your paper is 15 percent too long for the conference guidelines. You could submit the paper as is, and hope they take it anyway, but they might reject it. You could edit the paper to reduce the length, but you're tired of it by now, and you wish there was an easier way. So you shrink the font until the type fits on the page. Older guys like me can't read it without our glasses, but that's our problem, not yours.

In one review session, some supervisors simply used a smaller font to cram more than three accomplishments into the allocated space on the 333. Others went even further, jiggering the margins too, apparently on the theory that the blacker the page, the higher the preordered ranking would be.

After nominating myself as the font police for a year or two, I realized there was actually a bigger problem that we needed to confront directly: A performance review structure simply doesn't work well when run as laissez-faire capitalism.

NOT "EVERY MAN FOR HIMSELF"

The font changers were operating under a hidden assumption: that they were there to get their direct reports as high as possible on the final rank list. They believed that if all supervisors did the same thing, the result would be fair, similar to the US system of adversarial justice: "... the function of the courts is not to pursue the truth but to decide on the cases presented by the parties."

R&R sessions often operate under a quasi-adversarial philosophy, even though it can only be discerned from subtle clues. For instance, at one session, I felt that an important contribution made by someone else's employee had been overlooked. When I pointed this out, instead of being happy that an oversight had been avoided, the employee's supervisor reacted with suspi-

cion. What did I want in return for that favor?

Each supervisor participates in an R&R with substantial internal dissonance: Am I representing my direct report in the fairest possible way? Am I overlooking the importance of someone's work simply because I had less exposure to it? Am I factoring out personal likes or dislikes? If I think the group is making an error, am I intervening productively and effectively? Conversely, am I yielding to the common wisdom when that is the right thing to do? Opting not to go with the flow can be extremely difficult, especially when the technical leaders that the entire project depends on support the other point of view.

All of the supervisors participating in an R&R must have the attitude that their job is to help make the overall rank list as accurate as humanly possible. Supervisors do have a special responsibility to the direct reports they represent. But they also have a corporate responsibility. Over several years, these rank lists really do help shape the organization and its leadership, and they directly influence the company's fortunes.

Credit where credit is due

Many engineers are reluctant to brag. They feel that their accomplishments ought to speak for themselves. It can be difficult for these folks to strike the right balance between doing technical work as a member of a team and making sure that their own contributions are noticed and identifiable. We have all known engineers who went too far the other way, too: Every idea ever floated at a meeting became theirs. There is a happy medium.

The best recognition engineers get is from their peers, but the hierarchical management structure tends to dominate performance reviews. Intel has a process for collecting "360-degree feedback," gathering input from direct reports, peers, and managers. However, in my experience, it's very difficult to get objective peer feedback from this process. In addition, the information must be laundered so that the employee does not know who said what, which further diminishes the effectiveness of this source.

R&R sessions often operate under a quasi-adversarial philosophy.

I worked 12 months, you worked nine

The intended basis for Intel's technical meritocracy is an annual review of accomplishments—not "What are you good at?" "How well known are you?" "How many patents did you get two years ago?" "Hey, how about that Nobel Prize you won in 1994?"

Nice and simple, except for this: How do you handle people who were on sabbatical for a substantial portion of the year? What about someone who was sick for two months? Do you extrapolate the rest of the work they did, or do you stick to the actual accomplishments—in which case you may well find your top performer in an unaccustomed place, five slots down the rank group? And what about researchers whose work may not have visible payoffs every 12 months?

There are ways to handle these things, but it's not always easy to get them right.

"They" did this to you

Above all else, this syndrome must be avoided: "Well, Mike, I know you're unhappy that your ranking was in the third quartile. You should have been higher. I argued and argued for you, but the other supervisors just wouldn't listen."

An attempt at "you and me against the world" camaraderie is actually an admission of gross negligence on the supervisor's part. I have no patience at all for this ploy.

At the beginning of every rank session that I ran, I instructed the R&R managers that they were not to leave the room until they were satisfied that justice had been done to every person being ranked and rated, including their own direct reports.

Not managing to pay grade

Generally, the most awkward moments of an R&R come just after the participants have agreed to the ranking. This is because they have ranked the accomplishments with no consideration of an employee's grade level.

If the rank group has mixed grades—and I've seldom seen one that doesn't—it's not uncommon for some employees in lower pay grades to be rated higher than some other higher-grade employees. Indeed, these inversions are often an indicator that a promotion (or, rarely, a demotion) is appropriate.

The problem comes when the new rank list shows that a high-grade person is at the bottom of the list, indicating that the person is underproducing given his pay grade. This is bad news for the supervisor, who probably had considered this employee a successful member of the design project.

During the year, the supervisor may not have thought much about this person's pay grade—she just assigned tasks, and the employee completed them. Now the supervisor has to think back. Were the tasks grade-appropriate? Yes, the tasks were completed, but were they done on time? Is there an indication of declining productivity? Are there external factors that should be taken into account? For example, perhaps this person was newly promoted and is still catching up.

This is when the technical supervisors wish they had never agreed to do management.

When to promote

In addition to applying Intel's corporate policies about what fraction of a rank group should reasonably be promoted, the R&R managers also must try to identify employees who will remain successful at the next grade up.

Promoting people prematurely sets them up for failure, a failure that's not their own doing. But waiting until they have demonstrated their long-term ability to perform at the next level up doesn't always feel right either, especially when those people realize that their peers have been promoted but they have not.

Expectations, not quotas

In the same way that companies anticipate promotions, they also expect that some people at the bottom of the rank group might need to receive disciplinary action. The rank manager is held accountable for identifying these people, and it isn't easy.

No supervisor ever wants to appear at an R&R and say, "My employee isn't cutting it, and we need to formally begin the process that will either fix the problem or shove him out the door." That supervisor would feel like he's giving up on the employee and feel that he's failing as a manager.

Human nature being what it is, the supervisor will instead argue vigorously that there may be extenuating circumstances—transfer, divorce, sickness, personal relationships, or technical issues—and maybe something short of a formal proceeding should be done first. But if disciplinary action isn't taken when the situation warrants it, the organization's overall output could decline.

Establishing a quota of people to eliminate in every performance review period would be just plain wrong. If you consider that the organization is chopped up into 20-person rank groups of somewhat arbitrary membership, the top 20 performers of an entire 1,000-person organization could end up in one rank group. Insisting that the bottom 5 percent of every rank group should be disciplined would be ludicrous.

In my experience, Intel got this part of the performance review process just about right—enough pressure to make management take a serious look at stragglers, but no actual quotas.

Other companies that I have worked at conducted what they called performance reviews, but they weren't the same thing at all.

At the smaller companies, the supervisors just decided on their own who got what raises and promotions. One company had only one technical rank, so the only promotions were on the management side. These ad hoc performance review processes took much less management time, but they were ultimately ineffective.

Is Intel's performance review process perfect? Of course not. No review process is perfect. But Intel works hard on its review process and constantly tries to improve it. Intel should be proud of its efforts here. The elusive technical meritocracy is worth striving for. ■

Bob Colwell was Intel's chief IA32 architect through the Pentium II, III, and 4 microprocessors. He is now an independent consultant. Contact him at bob.colwell@attbi.com.

Not A Member Yet?

Do it now! Join the IEEE Computer Society today for many valuable benefits including...

Distance Learning

Members of the IEEE Computer Society enjoy FREE access to a comprehensive online training program for computer professionals. Join today and take any or all of 100 online course titles through our Distance Learning Campus. Subjects include Java, project management, Cisco networks, UNIX, Windows, HTML, Oracle, SQL, and more.

Magazines and Journals

To help you stay informed and competitive, the IEEE Computer Society publishes numerous specialized magazines and journals that are available to members at deeply discounted rates. Plus, you get a FREE subscription to *Computer* magazine included as part of your membership.

Conferences and Workshops

Enhance your knowledge and share practical experiences at more than 140 conferences, workshops, and symposia held each year worldwide. Members save at least 25% on registration fees and get advance notice of the meetings.

Other Valuable Benefits

- FREE membership in your local chapter
- FREE e-mail alias of Your.Name@computer.org
- FREE membership in up to four of 40+ Technical Committees
- FREE membership in 200+ Standards Working Groups
- Member discounts on hundreds of books and conference proceedings

IEEE COMPUTER SOCIETY

Join Today!

Complete the adjacent membership application and send it to the IEEE Computer Society office nearest you. For fastest service, apply online at http://computer.org/join

IEEE

2003 IEEE Computer Society Membership/Subscription Application

Membership and periodical subscriptions are annualized to and expire on 31 December 2003.
Pay full or half-year rate depending upon the date of receipt by the IEEE Computer Society as indicated below.

Membership Options*

All prices are quoted in U.S. dollars

	FULL YEAR Applications received 16 Aug 02 - 28 Feb 03	HALF YEAR Applications received 1 Mar 03 - 15 Aug 03
1 I do not belong to the IEEE, and I want to join **just** the Computer Society	$ 86 ❑	$44 ❑
2 I want to join **both** the Computer Society and the IEEE.		
I reside in the United States	$180 ❑	$91 ❑
I reside in Canada	$162 ❑	$82 ❑
I reside in Africa/Europe/Middle East	$158 ❑	$80 ❑
I reside in Latin America	$151 ❑	$76 ❑
I reside in Asia/Pacific	$152 ❑	$77 ❑
3 I already belong to the IEEE, and I want to join the Computer Society. (IEEE members need only furnish name, address, and IEEE number with payment.)	$ 37 ❑	$19 ❑

Are you now or were you ever a member of the IEEE? Yes ❑ No ❑ If yes, provide member number if known: _____

Special Interest Periodicals**

	ISSUES PER YEAR	FULL YEAR Applications received 16 Aug 02 - 28 Feb 03 PRINT	FULL YEAR ELECTRONIC	FULL YEAR COMBO	HALF YEAR Applications received 1 Mar 03 - 15 Aug 03 PRINT	HALF YEAR ELECTRONIC	HALF YEAR COMBO
IEEE Computer Society Digital Library	n/a	n/a	$109 ❑	n/a	n/a	$55 ❑	n/a
Computing in Science & Engineering	6	$41 ❑	$33 ❑	$53 ❑	$21 ❑	$17 ❑	$27 ❑
IEEE Computer Graphics & Applications	6	$38 ❑	$30 ❑	$49 ❑	$19 ❑	$15 ❑	$25 ❑
IEEE Design & Test of Computers	6	$36 ❑	$29 ❑	$47 ❑	$18 ❑	$15 ❑	$24 ❑
IEEE Intelligent Systems	6	$36 ❑	$29 ❑	$47 ❑	$18 ❑	$15 ❑	$24 ❑
IEEE Internet Computing	6	$38 ❑	$30 ❑	$49 ❑	$19 ❑	$15 ❑	$25 ❑
IT Professional	6	$39 ❑	$31 ❑	$51 ❑	$20 ❑	$16 ❑	$26 ❑
IEEE Micro	6	$36 ❑	$29 ❑	$47 ❑	$18 ❑	$15 ❑	$24 ❑
IEEE MultiMedia	4	$34 ❑	$27 ❑	$44 ❑	$17 ❑	$14 ❑	$22 ❑
IEEE Pervasive Computing	4	$40 ❑	$32 ❑	$52 ❑	$20 ❑	$16 ❑	$26 ❑
IEEE Security & Privacy NEW	6	$40 ❑	$32 ❑	$52 ❑	$20 ❑	$16 ❑	$26 ❑
IEEE Software	6	$43 ❑	$34 ❑	$56 ❑	$22 ❑	$17 ❑	$28 ❑
IEEE/ACM Transactions on Networking†	6	$40 ❑	$32 ❑	$48 ❑	$20 ❑	$16 ❑	$24 ❑
IEEE Transactions on:							
Computers	12	$40 ❑	$32 ❑	$52 ❑	$20 ❑	$16 ❑	$26 ❑
Knowledge & Data Engineering	6	$39 ❑	$31 ❑	$51 ❑	$20 ❑	$16 ❑	$26 ❑
Mobile Computing	4	$29 ❑	$23 ❑	$38 ❑	$15 ❑	$12 ❑	$19 ❑
Multimedia†	4	n/a	n/a	$32 ❑	n/a	n/a	n/a
Parallel & Distributed Systems	12	$39 ❑	$31 ❑	$51 ❑	$20 ❑	$16 ❑	$26 ❑
Pattern Analysis & Machine Intelligence	12	$42 ❑	$34 ❑	$55 ❑	$21 ❑	$17 ❑	$28 ❑
Software Engineering	12	$37 ❑	$30 ❑	$48 ❑	$19 ❑	$15 ❑	$24 ❑
Visualization & Computer Graphics	4	$31 ❑	$25 ❑	$40 ❑	$16 ❑	$13 ❑	$20 ❑
VLSI Systems†	6	n/a	n/a	$28 ❑	n/a	n/a	$14 ❑
IEEE Annals of the History of Computing	4	$30 ❑	$24 ❑	$39 ❑	$15 ❑	$12 ❑	$20 ❑

Payment Information

Payment required with application

Membership fee $_____

Periodicals total $_____

Applicable sales tax*** $_____

Total $_____

Enclosed: ❑ Check/Money Order****

Charge my: ❑ MasterCard ❑ Visa ❑ American Express ❑ Diner's Club

Card number _____

Expiration date (month/year) _____/_____ Signature _____

USA-only include 5-digit billing zip code ❑❑❑❑❑ IF3I

* Member dues include $16 for a 12-month subscription to *Computer*.

** Periodicals purchased at member prices are for the member's use only

*** Canadian residents add 15% HST or 7% GST to total. AL, AZ, CO, DC, NM, and WV add sales tax to all periodicals. GA, IN, KY, MA, MD, MO, NC, and OH add sales tax to print and combo periodicals. FL add sales tax to print and combo periodicals plus *Computer's* value. NY add sales tax to electronic periodicals.

**** Payable to the IEEE in U.S. dollars drawn on a U.S. bank account. Please include member name and number (if known) on your check.

† Not part of the IEEE Computer Society Digital Library. Electronic access is through www.ieee.org/ieeexplore.

For fastest service, apply online at http://computer.org/join

NOTE: In order for us to process your application, you must complete and return BOTH sides of this form to the office nearest you:

European Office
IEEE Computer Society
13, Avenue de l'Aquilon
B-1200 Brussels, Belgium
Phone +32 2 770 2198
Fax +32 2 770 8505
E-mail: euro.ofc@computer.org

Asia/Pacific Office
IEEE Computer Society
Watanabe Bldg.
1-4-2 Minami-Aoyama
Minato-ku, Tokyo 107-0062 Japan
Phone: +81 3 3408 3118
Fax: +81 3 3408 3553
E-mail: tokyo.ofc@computer.org

Publications Office
IEEE Computer Society
10662 Los Vaqueros Circle
PO Box 3014
Los Alamitos, CA 90720-1314 USA
Phone: +1 714 821 8380
Fax: +1 714 821 4641
E-mail: help@computer.org

Allow up to 8 weeks to complete application processing. Allow a minimum of 6 to 10 weeks for delivery of print periodicals.

Personal Information

Enter your name as you want it to appear on correspondence.
As a key identifier in our database, circle your last/surname.

Male ❑ Female ❑ Date of birth (Day/Month/Year) _____

Title First name Middle Last/Surname

Home address

City State/Province

Postal code Country

Home telephone Home facsimile

Preferred e-mail

Send mail to: ❑ Home address ❑ Business address

Educational Information

First professional degree completed Month/Year degree received

Program major/course of study

College/University State/Province Country

Highest technical degree received Program/Course of study

Month/Year received

College/University State/Province Country

Business/Professional Information

Title/Position

Years in current position Years of practice since graduation

Employer name Department/Division

Street address City State/Province

Postal code Country

Office phone Office facsimile

I hereby make application for Computer Society and/or IEEE membership and agree to be governed by IEEE's Constitution, Bylaws, Statements of Policies and Procedures, and Code of Ethics. I authorize release of information related to this application to determine my qualifications for membership.

Signature Date

APPLICATION MUST BE SIGNED

NOTE: In order for us to process your application, you must complete and return both sides of this form.

IEEE COMPUTER SOCIETY

IEEE

BPA Information

This information is used by society magazines to verify their annual circulation. Please refer to the audit codes and indicate your selections in the box provided.

A. Primary line of business

1. Computers
2. Computer peripheral equipment
3. Software
4. Office and business machines
5. Test, measurement and instrumentation equipment
6. Communications systems and equipment
7. Navigation and guidance systems and equipment
8. Consumer electronics/appliances
9. Industrial equipment, controls and systems
10. ICs and microprocessors
11. Semiconductors, components, sub-assemblies, materials and supplies
12. Aircraft, missiles, space and ground support equipment
13. Oceanography and support equipment
14. Medical electronic equipment
15. OEM incorporating electronics in their end product (not elsewhere classified)
16. Independent and university research, test and design laboratories and consultants (not connected with a manufacturing company)
17. Government agencies and armed forces
18. Companies using and/or incorporating any electronic products in their manufacturing, processing, research, or development activities
19. Telecommunications services, telephone (including cellular)
20. Broadcast services (TV, cable, radio)
21. Transportation services (airlines, railroads, etc.)
22. Computer and communications and data processing services
23. Power production, generation, transmission, and distribution
24. Other commercial users of electrical, electronic equipment and services (not elsewhere classified)
25. Distributor (reseller, wholesaler, retailer)
26. University, college/other education institutions, libraries
27. Retired
28. Others (allied to this field)

B. Principal job function

1. General and corporate management
2. Engineering management
3. Project engineering management
4. Research and development management
5. Design engineering management - analog
6. Design engineering management - digital
7. Research and development engineering
8. Design/development engineering - analog
9. Design/development engineering - digital
10. Hardware engineering
11. Software design/development
12. Computer science
13. Science/physics/mathematics
14. Engineering (not elsewhere classified)
15. Marketing/sales/purchasing
16. Consulting
17. Education/teaching
18. Retired
19. Other _____

C. Principal responsibility

1. Engineering or scientific management
2. Management other than engineering
3. Engineering design
4. Engineering
5. Software: science/management/engineering
6. Education/teaching
7. Consulting
8. Retired
9. Other _____

D. Title

1. Chairman of the Board/President/CEO
2. Owner/Partner
3. General Manager
4. V.P. Operations
5. V.P. Engineering/Director Engineering
6. Chief Engineer/Chief Scientist
7. Engineering Manager
8. Scientific Manager
9. Member of Technical Staff
10. Design Engineering Manager
11. Design Engineer
12. Hardware Engineer
13. Software Engineer
14. Computer Scientist
15. Dean/Professor/Instructor
16. Consultant
17. Retired
18. Other Professional/Technical

Developing the Distributed-Computing OS

Steven J. Vaughan-Nichols

Researchers are exploring the world of distributed computing, in which users in various locations can work with the same set of geographically dispersed resources. These efforts have led to such high-profile technologies as peer-to-peer (P2P), pervasive, and nomadic computing.

A critical part of this research is developing consistent approaches to distributed-computing operating environments, which must work consistently across many platforms and technologies. (See the sidebar, "Driving Forces behind the Distributed-OS Projects.")

Major efforts in this area include Globe, Opus, and Project Oxygen. None of the three represents radically new technologies, but instead each applies existing technologies in a novel way.

Globe (Global Object Based Environment), based at Holland's Vrije University, is developing a middleware network using distributed objects. Vrije computer science professor Marteen van Steen said, "The main significance of systems like Globe is that they will open up the road to the development of truly distributed Internet applications."

Opus, based at the US's Duke University, is a descendent of the University of California, Berkeley's WebOS project. Opus is designed to provide an overlay utility (used to manage a collection of nodes in a networking fabric) that dynamically allocates resources—CPU power, bandwidth, hard drive space—to requesting nodes in accordance with varying network characteristics and application demands.

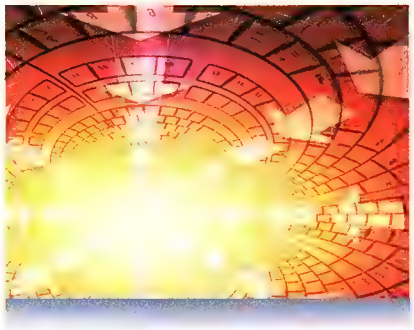

The Massachusetts Institute of Technology coordinates Project Oxygen, which is funded by the US Defense Advanced Research Projects Agency, Delta Electronics, Hewlett-Packard, Nokia, NTT, and Philips Electronics. Project Oxygen is an effort to combine wireless and other technologies to create a distributed network that makes personalized computing resources available to users no matter where they are in an environment.

Although not operating systems per se, these three projects attempt to abstract distributed and P2P computing into a single seamless operating environment where all resources on a distributed network can function like local resources.

GLOBE

The Globe project (http://www.cs.vu.nl/~steen/globe) is creating large distributed systems via distributed shared objects and their associated methods. Developers could generate applications using Globe as middleware instead of building network programs directly on top of the transport layer, as is the case today.

Active copies of objects, which communicate on a P2P basis, would be available simultaneously on every machine in a distributed system, and all users could invoke the objects' methods. A P2P approach would let systems run without a centralized repository of object types, thereby reducing network traffic and eliminating failures due to repository unavailability.

Globe improves a distributed system's functionality and speed by performing advanced tasks such as returning a Web page's content, accepting an e-mail message, accessing a file, or looking up a resources name in a directory.

Globe objects consist of five components: a control subobject, which manages client requests; a communications subobject, which enables interobject communications; a replication subobject, which manages the coherence of replicated objects; a security subobject, which controls access rights and object availability; and a semantics subobject, which implements the object's actions.

A developer would write the semantic subobject. The other subobjects then would either come from a library or be generated by the semantics object's requirements during compilation.

"The most important difference [between Globe objects and other objects] is that Globe offers support for objects that are physically distributed across multiple machines," van Steen explained. "Therefore, an object's state can be replicated and partitioned across multiple object servers."

Unlike other types of distributed

Driving Forces behind the Distributed-OS Projects

While academically driven, each of the three major distributed-OS approaches—Vrije University's Globe, Duke University's Opus, and MIT's Project Oxygen—addresses needs with commercial ramifications. These needs could become more pressing as distributed systems become more popular and as companies rely on them for more critical tasks.

The three projects are also striving for a unified, universal approach—rather than a proprietary one—to create a distributed-computing operating environment. Until now, various organizations have developed their own Web technologies to address different aspects of such environments, including Microsoft's DCOM (distributed component object model) and Corba's Internet inter-ORB (object request broker) protocol. Attempts to standardize this process have turned into political battles.

Globe's middleware approach would create an Internet-based distributed computing infrastructure that relies on a universal object format. Similarly, Project Oxygen is trying to provide a universal framework.

Another driving force behind the three projects is the desire to create a system in which computing resources throughout a network are available to users on demand, a goal of Opus. Industry is also pursuing this goal but on a proprietary basis.

objects, van Steen noted, each Globe object controls how its state is replicated, migrated, and otherwise spread across machines. In other words, Globe objects don't have to rely on another application or object request broker for these functions. Thus, a single Globe object model can provide a universal method to replicate and deliver information and service regardless of the underlying platform.

Because Globe objects provide the underlying network hooks, the technology would free Internet application developers from worrying about them. Rather than forcing developers to program fault tolerance, security, replication, and other functions into their applications, Globe acts as a middleware layer that provides these capabilities.

By providing distributed services, Globe objects could replace today's multiple, XML-based, Web services approaches and standards to deliver services over the Internet, said van Steen.

OPUS

Opus (http://www.cs.duke.edu/~dkostic/publications/opus-poster-sosp.pdf) is based on WebOS, a UC Berkeley project designed to provide OS services—including mechanisms for resource discovery and management, remote process execution, authentication, and security—to distributed applications.

After the Berkeley project ended in 1998, Duke University's Rebecca Braynard, Jeff Chase, Dejan Kostic, Adolfo Rodriguez, and Amin Vahdat used various WebOS concepts to develop Opus. The University of Texas at Austin and the University of Washington are also participating in the project.

Opus adds an overlay mechanism to the WebOS framework that lets applications transparently communicate their resource requirements to the underlying network and then use the provided resources.

This is important because on a single machine, application developers can rely on the local operating system to provide services. In a distributed system, however, application developers must build the services themselves from a hodgepodge of standards and multiple application servers, an approach that eats up programmer effort and system resources.

Opus solves this problem by providing the basic Internet-based OS services needed to build applications that are distributed, available, scalable, and dynamically reconfiguring.

The technology would thus provide a framework for delivering distributed P2P services, while giving Internet-based P2P applications much of the functionality generally found only in traditional client-server applications.

Opus researchers are developing a universal approach to accomplish this, rather than a custom overlay solution for a single application or network resource. The project's overlay approach would provide a common middleware interface for developers.

In Opus, the overlays serve as an API abstraction for accessing and using network resources, as determined by varying application demands and network characteristics such as available throughput. Opus lets an application access network resources even as the program's demands and network conditions change, a capability generally beyond today's distributed systems.

Opus would enable widespread access to Internet system resources. Vendors could use the technology commercially to deliver computing resources as a utility via the Internet to customers, said Amin Vahdat.

In the Opus architecture, shown in Figure 1, the overlay system works with the available network resources. A network administrator sets the basic resource allocation the system can use, while a network service provider agrees to provide minimum acceptable service levels.

In the topology-maintenance layer, the per-application overlay specifies the resources that a type of application needs, while the control overlay checks overlay nodes and maintains connectivity and performance characteristics among them.

Finally, the introspection layer checks network characteristics, examines the paths for accessing resources

to find the most efficient path, and determines which Opus overlay nodes (and accompanying functionality) should be allocated to a specific application.

PROJECT OXYGEN

Project Oxygen's (http://oxygen.lcs.mit.edu/) creators have an elaborate plan for developing a pervasive, distributed-computing approach that would make resources available when and where users need them, not just when users are near their own computers. The goal is to make computing part of the environment, rather than just a set of individual devices.

Nick Shelness, independent technology consultant and Lotus's former chief technology officer, said that Oxygen would make systems more user-centric. "That's very attractive and important," he explained.

Oxygen currently works with an infrastructure of handheld devices and workstations connected by IEEE 802.11b wireless LANs and Fast Ethernet networks.

The system uses Migrate, an MIT-developed host-mobility architecture designed to support mobile and intermittently connected network applications and to let legacy applications adapt to mobile environments. Users can migrate an active TCP connection across multiple IP addresses by sending new SYN packets (which establish virtual connections and synchronize packet sequence numbers at the start of a TCP connection) with the Migrate option enabled.

Oxygen's software underpinnings are mostly the result of MIT academic projects. For example, developers use IOA (input/output automaton), a programming language and set of tools designed for describing and building reliable distributed systems. IOA lets system developers express designs at different abstraction levels, from the high-level specification of global behavior to low-level versions that translate easily into code.

Project developers have implemented some Oxygen technology in their Intelligent Room, an interactive environment with embedded computation that participates in user activities. For example, microphone and camera arrays let the room "listen" to users, "observe" what they do, and transparently provide the information and communication resources they need.

Smart sketching and design tools let users express their ideas and let the room record them. Other tools enable collaborative work. An agent-based software infrastructure automatically handles many tasks that support the tools' operation.

Eventually, researchers hope to develop broad natural spoken and visual interfaces for accessing computing and communications resources.

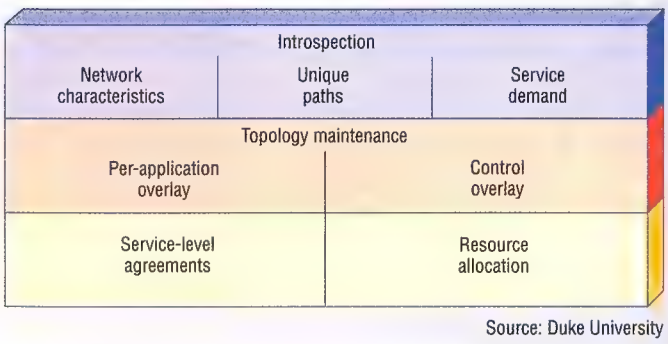

Figure 1. Opus architecture. Duke University's Opus adds an overlay mechanism to distributed systems that lets applications transparently communicate their resource requirements to the underlying network. The bottom layer describes the basic resource levels available to the system. The middle layer sets up and maintains the delivery of resources. The introspection layer determines which nodes should be allocated to an application and maintains performance levels.

Perhaps the greatest challenge the distributed-computing OS faces is the rise of Web services, represented by Microsoft's .NET and several Java-based approaches such as Sun Microsystems' Open Net Environment (ONE). Although no unifying organization for Web services exists yet, the technology is already popular in the enterprise.

Web services represent a set of platform-neutral technologies designed to ease the delivery of network services over intranets and the Internet. In essence, Web services integrate PCs, other devices, databases, and networks into a virtual computing fabric, which users work with via browsers. Clearly, this treads on much of the ground that distributed-computing operating systems hope to cover.

Dan Kusnetzky, vice president of system software research for market research firm IDC, said that while Globe, Opus, and Project Oxygen are interesting, they won't yield commercial products in the near future.

However, some distributed-OS proponents are more optimistic. For example, Victor Zue, director of MIT's Laboratory for Computer Science, said Project Oxygen could yield commercial products in the next three years.■

Steven J. Vaughan-Nichols is a freelance technology writer based in Arden, North Carolina. Contact him at sjvn@vna1.com.

Editor: Lee Garber, Computer, 10662 Los Vaqueros Circle, PO Box 3014, Los Alamitos, CA 90720-1314; l.garber@computer.org

Tiny "Punch Cards" Boost Storage Capacity

IBM researchers have created a postage-stamp-sized storage system with a density of 1 trillion bits per square inch, 20 times that of today's hard drives.

The new approach, part of IBM's Millipede project, stores data in a pattern of holes, a concept similar to that utilized by punch-card technology, which hasn't been popular since the 1970s. However, noted company researcher Peter Vettiger, IBM's system has some significant differences: the holes are only one-half-billionth of an inch wide, the equivalent of 10 nanometers or 50 atoms; and, unlike punch cards, the technology is rewriteable.

IBM anticipates Millipede storage will first be used in handheld devices because it offers a lot of memory in a small package, according to Vettiger. The Millipede chip is currently only 7 mm square, although it could be made bigger or smaller as necessary. "We could fit a Millipede device in a flash card format and offer at least five to 10 times more storage," noted Vettiger.

To write data, Millipede uses a tiny silicon tip heated to 400°C to dent a thin polymer layer on a silicon chip. The research team has been experimenting to find the polymer that offers the least mechanical wear and most stability.

To read data, the polymer layer passes under a silicon tip heated to 300°C. When the tip hits a dent, its temperature drops because more of its surface is in contact with the cooler polymer. The lower temperature drops the tip's electrical resistance, which can then be measured to read the indentation as a one or a zero. To erase data, a hot silicon tip closes the indentations.

Reading and writing data using a single silicon tip currently take 1,000 times as long as a hard-drive system. Millipede researchers have accelerated the process via a prototype with 1,024 silicon tips working in parallel. And because using more tips minimizes the distance the polymer layer must move, the practice also reduces power consumption. IBM is currently working on a prototype with 4,000 tips.

According to IBM, the Millipede chip could be made using existing techniques with few modifications, which means the manufacturing process would be relatively inexpensive.

Millipede addresses problems caused by the way other types of storage devices store data. For example, electronic devices store data as groups of electrons. As the devices get smaller, so do the electron groups. However, Vettiger explained, this results in a lower stored electrical charge to hold the electrons together, which could cause the electron groups to fall apart and the data they store to be lost. Researchers are working to solve this problem.

Making magnetic storage smaller can also affect data quality, Vettiger added. Temperature changes more strongly affect smaller magnetic particles, he said, causing them to move around and lose the orientation that establishes whether they represent a one or a zero.

"It's still very much at the laboratory level, but they have proof of concept," said Marlene Bourne, a senior analyst at the In-Stat/MDR market research firm. "From a commercial standpoint, it's hard to say when and if this will come on the market. Other companies are working on similar concepts, so it seems they are on to something." ■

—Linda Dailey Paulson

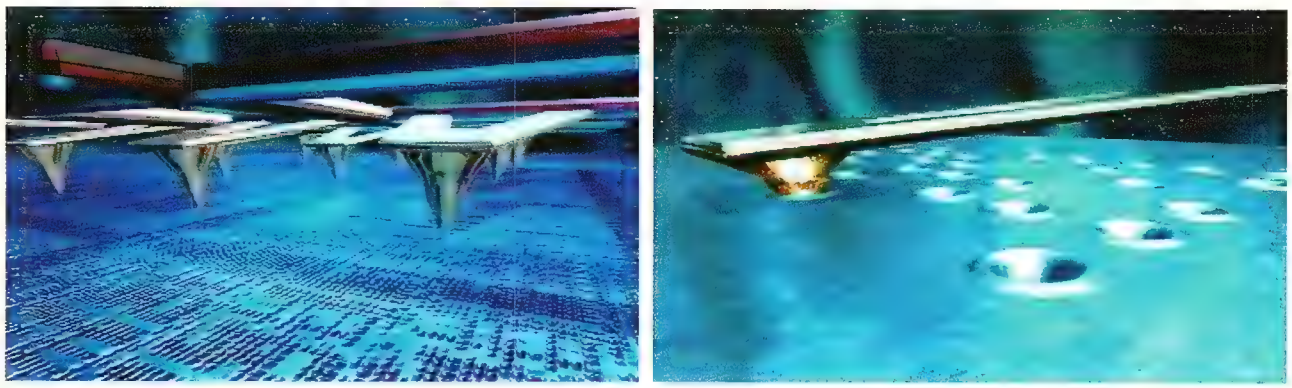

IBM's high-density storage system uses tiny, heated silicon tips to make dents in and thereby write data to a polymer layer on a silicon chip. The tips can also be used to read or erase data.

Researchers Bring Wireless Communications to the Chip

A team of researchers has built a wireless communications system entirely on a chip. This system could greatly improve communications within a chip. Proponents also see the technology as being useful for chip-to-chip communications and stand-alone radio applications.

The University of Florida research team, headed by Professor Kenneth Kyongyop O, placed a tiny radio transmitter and antenna on a common 0.18-micron CMOS chip. The 15-GHz antenna is just 2 millimeters by 10 microns. The receiver circuit measures 400 by 500 microns. In the near future, O said, he hopes to be using an even smaller, more powerful antenna.

The technology's ability to improve intrachip communications becomes particularly important as chips get bigger, said Peter Glaskowsky, editor in chief of *Microprocessor Report*, a technical publication for the microprocessor industry.

Currently, the wireless system would be used to transmit clock signals, which are sent to all parts of a processor simultaneously to synchronize the chip's various information-processing tasks. Sending information to all parts of a processor at the same time becomes more difficult as on-chip distances increase. The challenges are timing accuracy and signal distortion over longer distances, O noted. Transmitting signals wirelessly, rather than over varying lengths of wire, eliminates this problem.

"Our technology is ideally suited for the clock signal, but it could be used for other global signals such as when the chip is put into sleep mode," O said. The wireless systems could also be used to replace other types of on-chip wiring, provide processor-to-processor communications, and put microphones or receivers on chips for use in rescue or eavesdropping operations.

However, said Glaskowsky, the technology does not have enough bandwidth to replace many existing wired connections between devices, such as PCI.

The University of Florida research team's primary development challenge was avoiding interference in data transmissions across a chip, a problem compounded by microprocessor noise. "One of the things that is helping us," O explained, "is we are operating [our] circuit at a higher frequency than the rest of the circuits."

Nonetheless, Glaskowsky noted, using wireless approaches will make the chip-manufacturing process longer and more expensive.

O's research is sponsored primarily by the Semiconductor Research Corp., an industry consortium. He said his team may develop sufficient data to attract commercial development or applications within the next two years. ■

—Linda Dailey Paulson

Editor: Lee Garber, Computer, 10662 Los Vaqueros Circle, PO Box 3014, Los Alamitos, CA 90720-1314; l.garber@computer.org

War Chalking for Free Wireless Networking

Years ago, homeless people wrote symbols on posts or other surfaces to help one another identify locations where they might find free food, avoid a person who might threaten their safety, or provide basic survival information.

Today, some travelers using wireless technology are beginning to undertake a similar activity. As they move about, these travelers exploit the IEEE 802.11 wireless-LAN technology's broadcast range to access nearby users' Internet service, sometimes without the hosts' permission, on a laptop or handheld device.

In some cases, travelers put chalk marks on sidewalks or other surfaces to identify the status of an IEEE 802.11 network they have discovered, a practice called war chalking. The name is based on war dialing, the term hackers use when they dial many phone numbers trying to find available modems, and also war driving, the term for traveling around looking for a WLAN access point.

War chalking has been discussed for a while within the IEEE 802.11 user community, noted Tim Pozar, founder of the San Francisco-based Bay Area Wireless Users Group, in which paying WLAN customers use the technology's transmission range to provide their neighborhoods with free Internet service.

Some people are actively promoting war chalking, which now even has a Web site (http://www.warchalking.org). Nonetheless, Pozar said the practice would not be useful for many people.

If war chalking catches on, Internet service providers won't be happy. ISPs already maintain that war driving unfairly deprives them of customers who otherwise would have to pay for Internet services. ■

—Linda Dailey Paulson

New Chips Will Help Computers Visualize in Three Dimensions

Two companies have taken innovative approaches to developing chips that let computers visualize objects in three dimensions. This could improve how computer-based systems interact with objects and could be used in applications such as surveillance or robotic navigation.

Tyzx Inc.'s DeepSea chip, according to president and CEO Ron Buck, develops a 3D vision of an object by using stereo vision—determining how far an object or parts of an object are by comparing the slightly different views from each of two cameras that serve as imagers.

The low-power, 33-MHz DeepSea chip uses an algorithm called Census that quickly finds corresponding left and right images in the two imagers' streams. The host system then compares the pixels in the similar images, a computing-intensive task, and calculates the distance.

Tyzx said most of its first customers have been research laboratories. For example, MD Robotics, an aerospace company with research facilities, is working with DeepSea technology to build devices that would retrieve satellites in orbit.

The technology could also be used in the military in unmanned ground vehicles or in robotics used to explore underwater or uninhabitable areas.

Canesta Inc. has used a different approach in its chip that can recognize 3D images. A host computer system emits infrared light that bounces off of an object. The chip measures the time it takes for the light to bounce back. The software then determines the distance.

The chip, which Canesta hopes to release commercially later this year, could be used in such applications as car-based occupant-sensing technology to help in air-bag deployment, said Jim Spare, the company's vice president of product marketing.

Researchers have been working on 3D-visualization computer technology since the 1960s, noted Daniel Scharstein, assistant professor at Middlebury College's Department of Mathematics and Computer Science. Companies such as Tyzx are moving forward, he said, but "it will be at least 10 to 20 years before it becomes mainstream. Other breakthroughs are needed to make it robust." ■

—Linda Dailey Paulson

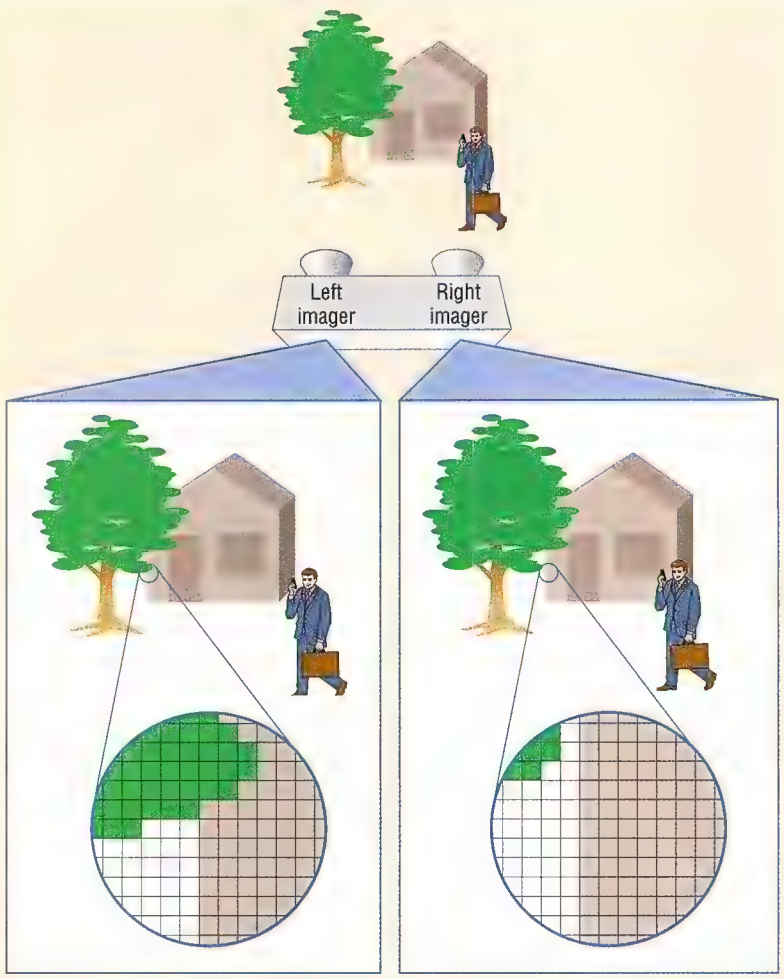

Source: Tyzx Inc.

Tyzx Inc.'s DeepSea system uses stereo-vision technology to enable computers to visualize objects in three dimensions. Two cameras capture images from different angles. The DeepSea chip then uses an algorithm that quickly finds corresponding left and right images in the two imagers' streams. The host system then compares the pixels in the similar images and calculates the object's distance.

Merging Fossil Specimens with Computer-Generated Information

Oliver Bimber
Bauhaus University Weimar
Stephen M. Gatesy
Brown University
Lawrence M. Witmer
Ohio University
Ramesh Raskar
Mitsubishi Electric Research Laboratories
L. Miguel Encarnação
Fraunhofer Center for Research in Computer Graphics

Augmented paleontology—the use of augmented-reality technologies to clothe fossil skeletons with soft tissues and skin—will let paleontologists bring their bare-bones specimens to life.

Paleontology is filled with mysteries about the plants and animals that lived thousands, millions, even billions of years before the first humans walked the earth. To answer questions about these organisms, paleontologists rely on the excavation, analysis, and interpretation of fossils. Embedded and preserved in the earth's crust, fossils are the remains or traces of ancient life forms, including bones, teeth, shells, leaf imprints, nests, and footprints.

Fossils can disclose how organisms evolved over time and their relationship to one another. While they reveal much, such as the general shape and size of ancient living things, fossils keep us guessing about these organisms' color, sound, and—most significantly—their behavior.

For several years, modern paleontologists have used 3D computer graphics to help reconstruct these pieces of the past.[1] State-of-the-art scanning technology produces 3D fossil replicas that scientists can process and study without physical constraints.

Paleontologists typically generate volumetric data sets for analysis, such as magnetic resonance imaging or computed axial tomography scans, and they use surface models for digital preservation and reproduction. To study ontogeny—an organism's growth and form—paleontologists apply mathematical models for simulation and visualization. Likewise, computer animations help study dinosaur locomotion. Beyond building knowledge of our world, the results of this work influence how dinosaurs appear in museums, illustrations, and movies, and as toys.

In the past 40 years, technological advances have continued to blur the boundary between real and computer-generated worlds. Augmented reality leverages this technology to provide an interface that enhances the real world with synthetic supplements. Paleontologists can use AR to present virtual data, such as 3D computer graphics, directly within a real environment rather than on a flat monitor. We coined the term *augmented paleontology* to refer to the application of AR to paleontology. AP seeks to support paleontologists in their research, and communicate the results of paleontology to museum visitors in an exciting and effective way.

0018-9162/02/$17.00 © 2002 IEEE

Augmented-Paleontology Tools

Paleontologists can use the following tools, which embody several developing technologies, to create augmented paleontology displays.

The Virtual Showcase

This museum display provides an imaginative method for accessing, presenting, and interacting with scientific and cultural content. Conceptually, the Virtual Showcase is compatible with conventional showcases. It also allows the display of computer-generated 3D graphics, animations, and real artifacts within the same 3D space.

Potentially, such interactive presentations can be more entertaining and engaging than conventional, passive showcases. The Virtual Showcase turns the exploration of cultural and scientific objects into an interactive process that can enhance the museum visitor's experience and facilitate the learning process.

Driven by off-the-shelf PCs with conventional 3D graphics cards, the Virtual Showcase, shown in Figure A, provides stereoscopic viewing, a high and scalable resolution, better support for eye accommodation, mutual occlusion between real and virtual objects, and multiple user support. The Virtual Showcase is a projection-based, optical see-through display that consists of two main components: a convex assembly of half-silvered mirrors and a graphics display. It uses video projectors as controllable light sources that illuminate the showcase's contents on a per-pixel basis while the system presents view-dependent stereoscopic graphics to observers. A wireless infrared tracking device provides user tracking, which lets the system render the augmented graphics from the correct perspective.

Figure A. Virtual Showcase. This example of a cone-shaped prototype supports two to three users, wireless user tracking, and a seamless surround view.

Video mixing

Originally developed for television special effects, video mixing combines video streams from cameras with computer-generated graphics, merging captured images of the real environment with synthetic images prior to displaying or recording them.

Using this technique to augment a 3D environment with computer graphics requires continuous knowledge of the physical cameras' parameters to ensure consistency between the two streams. To realize this goal, researchers apply computer vision algorithms to analyze captured video images and detect integrated landmarks within the real environment. These algorithms use knowledge of the landmarks' properties and of the camera's internal parameters, such as field of view and focal length, to reconstruct the camera's external parameters—its position and orientation. Knowing this information, a corresponding virtual camera can be defined that will render the graphical augmentation from the physical camera's perspective.

Projector-based augmentation

By replacing a physical object—with its inherent color, texture, and material properties—with a neutral object and projected imagery, projector-based augmentation can directly reproduce either the object's original or altered appearance. This approach effectively lifts the object's visual properties into the video projector.

Projector-based augmentation is ideal when the physical object is available for visualization, even if it forms a complex geometric shape. Multiple users can view the augmented object simultaneously, without using stereo glasses or head-mounted tracking devices.

For a high-quality augmentation, the system computes an image of the original object's textured 3D graphics model from the projector's viewpoint. When projected, the rendered image appears smoothly registered with the neutrally colored object, which changes its appearance accordingly.

An interdisciplinary team of paleontologists, graphics designers, and computer scientists has already applied the AP interface to soft-tissue reconstruction and the study of dinosaur locomotion.

SOFT-TISSUE RECONSTRUCTION

Despite the volumes of data paleontologists have already amassed, many questions about dinosaurs remain. Paleontologists seek to discover what dinosaurs looked like and how they breathed, smelled, and ate. Generally, only the fossilized bones and teeth—the hard parts—of dinosaurs are preserved. But these creatures' soft tissues, which formed their bodies and animated their bones, provide the keys to unlocking the secrets of dinosaur biology.

Previous research sought to remedy this situation by reconstructing soft-tissue components such as muscles, veins, arteries, and cartilage.[2] Soft tissues carve distinct marks in the bone's structure that paleontologists can compare to the marks found on the bones of modern-day animals.

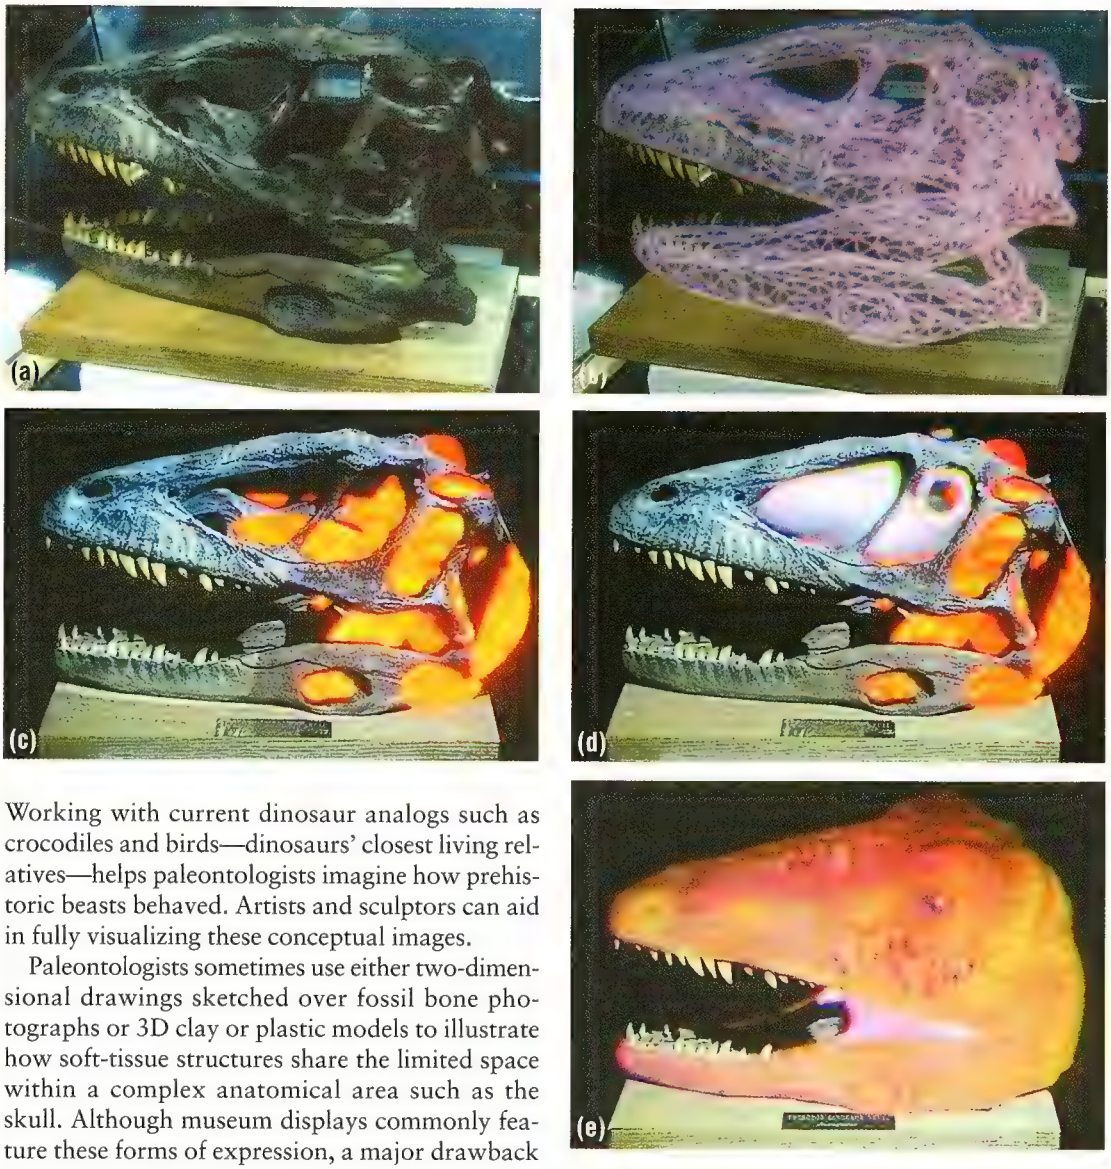

Figure 1. Visualization with the Virtual Showcase: (a) the physical skull of Deinonychus is placed inside the display, (b) a scanned skull geometry is registered to the real counterpart, (c) different muscle groups are augmented, (d) the paranasal air sinuses and the bony eye rings are integrated into the skull, and (e) the skin is superimposed on the skull.

Working with current dinosaur analogs such as crocodiles and birds—dinosaurs' closest living relatives—helps paleontologists imagine how prehistoric beasts behaved. Artists and sculptors can aid in fully visualizing these conceptual images.

Paleontologists sometimes use either two-dimensional drawings sketched over fossil bone photographs or 3D clay or plastic models to illustrate how soft-tissue structures share the limited space within a complex anatomical area such as the skull. Although museum displays commonly feature these forms of expression, a major drawback is that they lack dynamics and interactivity. Further, modifying these static drawings and sculptures is difficult.

Paleontologists can use AP to dynamically visualize reconstructed soft tissues that have been fully integrated with the fossilized bone structure. Then they can assess the conformational relationships of the reconstructed components to test the soft-tissue anatomy model. For example, paleontologists can use AP to directly assess how diverse anatomical elements, ranging from the eyeballs and tongue to the jaw muscles and nasal cartilage, share the skull's limited interior. AP lets paleontologists study how soft tissues accomplish biological tasks such as muscle contraction of the jaw and how the bulging of those contracting muscles dynamically affects the conformation of surrounding structures.

For museum visitors, AP combined with digital storytelling enables a more exciting and interactive experience, and it has the potential for improving knowledge transfer.

FROM ANATOMICAL PACKING TO DIGITAL STORYTELLING

To demonstrate AP's benefits, consider the cast skull of a Deinonychus, an Early Cretaceous dinosaur, augmented with 3D computer models of reconstructed soft tissues and missing bones. Observers can perceive the real and virtual components together in stereo from any perspective using the Virtual Showcase,[3] a special projection-based AR display device that the "Augmented-Paleontology Tools" sidebar describes in detail.

The bone and soft-tissue data sets for Deinonychus have been acquired from different sources. The skull bones, for example, were captured using 3D laser-scanning technology. After paleontologists assembled reconstructed muscles, eyeballs, nostrils, ears, sinuses, and skin in the lab, media designers premodeled them with an off-the-shelf modeling tool.

As Figure 1 shows, to visualize the packing process and present the final results with Virtual Showcase,

Figure 2. Superimposing foot motion: A mix of rendered graphics and live video superimposes the animated theropod foot skeleton over the real track.

the physical skull is placed inside the 3D display, then augmented with the reconstructed soft tissues.

The scanned geometric representation must be registered against its physical counterpart first. Doing so lets the Virtual Showcase compute the illumination and occlusion effects directly onto the skull's surface. A simple mouse-based interface can then interactively place the premodeled soft tissues inside the skull so that the paleontologists can investigate contact points and collisions with the physical bones and other soft tissues.

The display uses a stereoscopic graphics presentation to position the virtual components within the same 3D space as the skull bones. Multiple users can wear head-tracking devices to walk around the display and simultaneously observe the augmented artifact from different perspectives.

Simulating realistic occlusion effects between the physical bones and the virtual soft tissues is essential for packing and presentation. The Virtual Showcase is an *optical see-through display*. A drawback of such displays is that achieving realistic occlusion effects between real and virtual components is difficult. To overcome this difficulty, the Virtual Showcase uses controllable video projectors instead of simple lightbulbs.

The video projectors create view-dependent lighting effects on the real skull's surface. Generating shadows on the physical object exactly beneath the overlaid graphics, for example, lets virtual parts mutually occlude the underlying real surfaces. Having the skull's depth information, on the other hand, lets us cull the occluded graphics before they display. Thus, the physical bones can occlude virtual components and vice versa. This strongly enhances the interactive packing process and the presentation's realism.

To simulate interaction and behavior in different situations, we animate the virtual components during packing and presentation. To render the graphics, we use a conventional game engine that provides both high-quality animation and interactive frame rates. Enhancing the presentation with synchronized audio output and projector-based illumination lets us achieve an effective form of digital storytelling by dynamically fading in and out specific parts of the physical skull. During the presentation, different soft-tissue layers and components display over time, while a variety of multimedia aids—such as voice, text annotations, graphical animation, and lighting effects—explain their functions and relationships.

STUDYING LOCOMOTION

Paleontologists use AP to analyze fossilized dinosaur footprints left by theropods 210 million years ago.[4] Discovered on rocky exposures in eastern Greenland, the tracks of these bipedal carnivorous dinosaurs reveal how they moved about on two legs and how their locomotor pattern evolved over time.

These footprints can help paleontologists discern the similarities and differences between early theropods and birds, their living descendants. Shallow footprints made on firm mud record the shape of the bottom of the foot, but provide little information about how the limb moved while on the ground. Many of the Greenlandic trackways, however, were made by theropods that sank to varying depths in soft mud. Such deep prints record the path of the foot through a volume of sediment in three dimensions, thereby allowing reconstruction of limb motion.

This dynamic perspective has helped yield functional explanations of the deep tracks' many unusual

features. For example, the elongate front of deep Greenlandic tracks strongly resembles the tracks living birds make, as when turkeys walk through deep mud. This finding verifies that the early theropods' toes converged as they lifted them from the substrate, a feature that many birds still retain today.

Other features provide evidence of important differences. Birds leave no sole prints because they quickly lift their ankles after contacting the ground. In contrast, all deep tracks from early theropods show a substantial sole print, indicating that the ankle did not rise up until much later in the stride cycle. The impression left by the first toe points backward in Greenlandic tracks, whereas the toe itself points forward in all of this age's fossil skeletons. Although characteristic of perching birds, a reversed first toe is surprising in these ancient theropods. In this case, however, the fossil record may be misleading. Three-dimensional computer simulations that use a particle system reveal that a forward-pointing first toe *can* create a backward-pointing slash in the surface as it plunges down and forward into the mud.

Tracking the motion

Video-based AR may help paleontologists analyze limb motion in extinct animals by allowing animated models to interact with physical casts of dinosaur footprints. To capture the live video stream of the surrounding environment, we use a desktop-based video mixing configuration and a conventional Webcam. The system uses a pattern-matching algorithm to recognize integrated markers, detect their appearance within the images, and estimate the camera's position and orientation relative to the markers. We then apply these parameters to set up a corresponding virtual camera that renders the graphics scene from the same perspective as the real camera. Finally, we merge both images and display them in real time on a desktop screen. As Figure 2 shows, this technique lets us superimpose a computer-animated model of a theropod's foot skeleton over a real track and verify the animation with the physical toe marks on the imprint.

Moving into the spotlight

In many cases, original fossils are not available to either museums or paleontologists. Further, local politics or concerns that moving the fossil might damage or destroy it can limit its availability. Even when fossils are available, various researchers and institutions often must share them.

In these situations, handmade casts sometimes must substitute for the real fossils. A cast can easily express the fossil's shape, but not its original texture and detail information, such as skin impressions. However, photographs *can* capture this information.

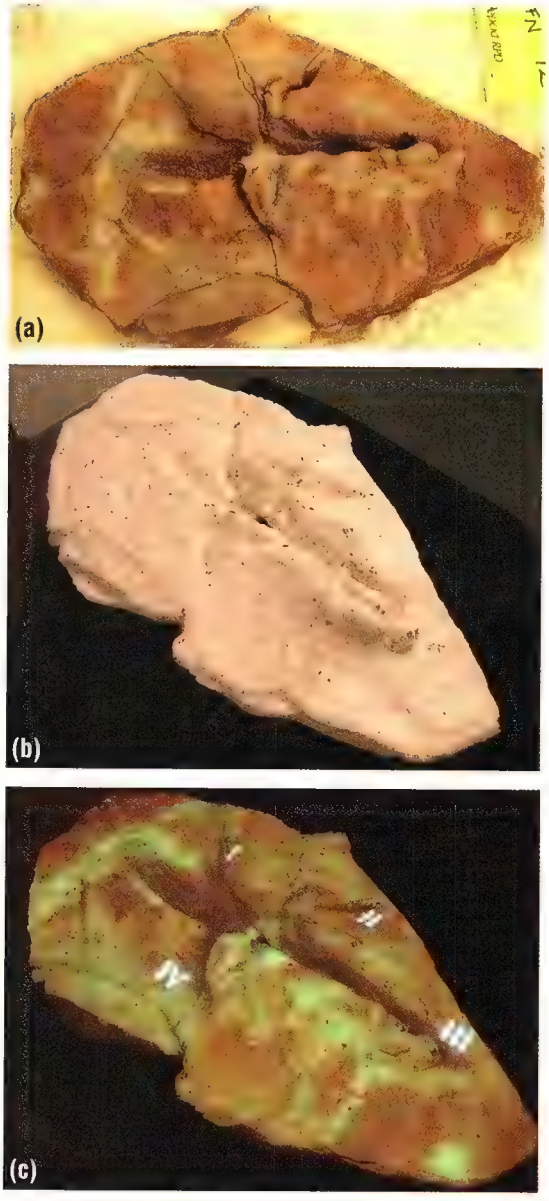

Figure 3. Projector-based augmentation of a therapod track. (a) The original track, preserved in Greenland's Triassic Fleming Fjord formation; (b) handmade cast of the shallow track; and (c) texture of the original track, with additional annotations that indicate the imprints of the toes, projected directly onto the cast.

Instead of painting the cast with a texture, we use video projectors to seamlessly map existing photographs and rendered images directly onto the cast's surface. Figure 3 shows how combining a projector-based augmentation[5] with physical replicas can recreate artifacts in laboratories and museums located hundreds or thousands of miles from the original specimen.

A projector-based augmentation also provides interactivity. The display can dynamically change the lighting conditions and the surfaces' color properties. The augmentation can display additional information, such as annotations and highlights, directly on the cast as part of an interactive or linear storytelling installation. In contrast to a completely virtual model such as a textured 3D scan, a

physical cast provides haptic and tactile feedback and natural, autostereoscopic, 3D perception.

The past decade's technological progress has improved hardware and software significantly, opening potential application areas for AR in general[6] and AP specifically. Although today's AR technology resembles the early stages of virtual reality, researchers foresee AR developing at a much faster rate and offering a wider range of applications.

Initially, researchers will use AR for applications that do not require high-precision technology. Within the paleontology domain, AR should soon be feasible for educational purposes. Interactive digital storytelling setups in museums, for example, would let paleontologists display and communicate their findings more effectively. Such installations must, however, be stable, childproof, and affordable.

In terms of supporting paleontologists in their research, technology must evolve and become more cost effective. Specifically, researchers must improve displays and tracking technology and make them affordable so that paleontologists, preparators, and restorators can effectively apply AP in their laboratories.

In the long term, mobile AR may become robust enough to support paleontologists during field trips and at dig sites. However, outdoor AR presents a much greater challenge than indoor AR, given the associated environment's larger scale and limited adaptability. User tracking and controlled illumination remain two of the main challenges confronting mobile AR. ■

ACKNOWLEDGMENTS

This work is supported in part by EU IST-2001-28610 to OB, NSF DBI-9974424 to SMG, and NSF IBN-9601174 to LMW. LMW lab's Casey Monahan Holliday assisted with Deinonychus's soft-tissue reconstruction. Julia Wolin and Milagros Leon of Rhode Island School of Design/CRCG created multimedia content and animations.

REFERENCES

1. S. Figgins, "Graphics in Computational Paleontology," *IEEE Computer Graphics & Applications*, vol. 21, no. 6, 2001, pp. 6-13.
2. L.M. Witmer, "The Evolution of the Antorbital Cavity of Archosaurs: A Study in Soft-Tissue Reconstruction in the Fossil Record with an Analysis of the Function of Pneumaticity," *J. Vertebrate Paleontology*, vol. 17, supplement to no. 3, 1997, pp. 1-73.
3. O. Bimber et al., "The Virtual Showcase," *IEEE Computer Graphics & Applications*, vol. 21, no. 6, 2001, pp. 48-55.
4. S.M. Gatesy et al., "Three-Dimensional Preservation of Foot Movement in Triassic Theropod Dinosaurs," *Nature*, vol. 399, 1999, pp. 141-144.
5. R. Raskar et al., "Shader Lamps: Animating Real Objects with Image-Based Illumination," *Proc. Eurographics Rendering Workshop* (EGRW 2001), Springer-Verlag, New York, 2001, pp. 89-102.
6. R. Azuma et al., "Recent Advances in Augmented Reality," *IEEE Computer Graphics & Applications*, vol. 21, no.6, 2001, pp. 34-47.

Oliver Bimber *is a scientist on the media faculty at the Bauhaus University in Weimar, Germany, and a former senior researcher at the Fraunhofer Center for Research in Computer Graphics in Providence, R.I. He initiated the Virtual Showcase project in Europe and the Augmented Paleontology project in the US. Bimber is a PhD candidate at the Technical University of Darmstadt. Contact him at oliver.bimber@medien.uni-weimar.de.*

Stephen M. Gatesy *is an associate professor of biology in the Department of Ecology and Evolutionary Biology at Brown University, Providence, R.I., where he works as an evolutionary functional morphologist. Gatesy received a PhD in organismic and evolutionary biology from Harvard University. Contact him at stephen_gatesy@brown.edu.*

Lawrence M. Witmer *is an associate professor of anatomy at the Ohio University College of Osteopathic Medicine, Athens, Ohio. Witmer received a PhD in anatomy from the Johns Hopkins Medical School. Contact him at witmerl@ohio.edu.*

Ramesh Raskar *is a research scientist at MERL in Cambridge, Mass. He initiated the Shader Lamps project at the University of North Carolina's Computer Graphics Lab. Raskar received a PhD in computer science from the University of North Carolina, Chapel Hill. Contact him at raskar@merl.com.*

L. Miguel Encarnação *is the head of the Human Media Technologies Department at the Fraunhofer Center for Research in Computer Graphics and an adjunct professor of computer science at the University of Rhode Island. Encarnação received a PhD in computer science from the University of Tübingen, Germany. Contact him at me@crcg.edu.*

The Importance of Branching Models in SCM

Among the branching models used in software configuration management, the branch-by-purpose model offers better support for parallel development efforts and improved control of both planned and emergency software releases.

Chuck Walrad
Davenport Consulting

Darrel Strom
Expert Support

If you want to improve software quality, you must first understand your software. What are its pieces? How are they organized and related to one another? If you do not understand your code base, your odds of updating it without breaking something are poor.

All too often, we see projects brought to their knees while trying to get the software out the door to testing groups or to customers for acceptance testing. The developers have worked feverishly to get the features in, and the testers are waiting to test, but the process fails at the integration point. The software components don't hang together, and some components may be missing. Wrong versions get distributed, and previously fixed bugs somehow reappear.

Why don't organizations have a better handle on their software? Is the problem a function of code size and complexity, is it inherent in parallel development efforts, or is it simply a result of staff turnover or cutting corners to meet schedule pressures?

All of these factors can contribute to the situation, but the real problem lies in a fundamental misunderstanding of software configuration management as it applies to real-world application development.

SCM DEFINED

Software configuration management serves two different functions:

- Management support for controlling changes to software products. This function includes the activities classically associated with SCM[1-3]—specifically, identifying the software components, controlling changes to them, recording and reporting component and configuration status, and conducting audits and reviews.
- Development support for coordinating file changes among product developers.[4,5] These activities include file version identification, software building, and release management.

Branching is integral to version management, software build correctness, and release management. It enables parallel development of a new system and provides concurrent support of multiple releases by labeling each instance of a branched configuration item and establishing a mapping between the label and the module revisions, as described in the "SCM Glossary" sidebar.

Good decisions about when and why to branch can make it much easier for developers and release engineers to coordinate software product changes. The right branching strategy makes it easier to deliver the right code, re-create past releases, and—if necessary—roll back to a previous release.

Adopting the right SCM branching model facilitates rapid development, increases overall product quality and process efficiency, reduces the incidence of software failures, and improves organizational performance.

THE BRANCHING MODEL

A branching model embodies the rationale

0018-9162/02/$17.00 © 2002 IEEE

SCM Glossary

Some basic definitions are helpful in developing a fundamental understanding of software configuration management in the real world of application development.

Baseline. In software development, the IEEE standards define a baseline as a "specification or product that has been formally reviewed and agreed upon, that serves thereafter as the basis for further development, and that can be changed only through formal change control procedures."[1] Alternatively, a baseline can be described as a "set of software items formally designated and fixed at a specific time during the software life cycle," or it can "refer to a particular version of a software item that has been agreed upon. In either case, the baseline can only be changed through formal change control procedures."[2]

Branch. "A branch is an agreed upon split of an item [item, product, or system] into multiple iterations [identifying each] ...instance of item, product, or system, ...[providing] an exact mapping between a version label and module revisions."[3]

Configuration management or software configuration management. (1) "SCM involves identifying the configuration of the software at given points in time, systematically controlling changes to the configuration, and maintaining the integrity and traceability of the configuration throughout the software life cycle. The work products placed under SCM include software products that are delivered to the customer and the items that are identified with or required to create these software products."[4] This includes the tool chain used to create, test, and maintain the product. (2) "A discipline applying technical and administrative direction and surveillance to: identify and document the functional and physical characteristics of a configuration item, control changes to those characteristics, record and report change processing and implementation status, and verify compliance with specified requirements."[1]

Release. The distribution of a software configuration item outside the development activity. This includes internal releases as well as distribution to customers.[5]

Release engineering. The process of moving a configuration through the software life cycle and delivering the finished product in the appropriate format and media.

Release management. Identification, packaging, and delivery of a product's elements, such as the executable, documentation, release notes, and configuration data.[5]

Software configuration item. An SCI aggregates software designated for configuration management and treats it as a single entity in the SCM process.[1]

Version. (1) "An initial release or re-release of a software configuration item, associated with a complete compilation or recompilation of the SW configuration item."[1] (2) "An initial release or complete re-release of a document, as opposed to a revision resulting from issuing change pages to a previous release."[1] (3) "A particular identified and specified software item."[5]

References

1. *IEEE Std. 610.12-1990, Standard Glossary of Software Engineering Terminology*, IEEE Press, Piscataway, N.J., 1990.
2. *Software Engineering Body of Knowledge*, trial version, IEEE Press, Piscataway, N.J., 2001, p. 108.
3. M. Ben-Menachem, *Software Configuration Guidebook*, McGraw Hill, Maidenhead, Berkshire, UK, 1994.
4. M.Paulk, et al., *Key Practices of the Capability Maturity Model, Version 1.1*, SEI-93-TR-25, Software Eng. Inst., Carnegie Mellon University, Pittsburgh, 1993.
5. *Software Engineering Body of Knowledge*, trial version, IEEE Press, Piscataway, N.J., 2001, p. 111.

adopted for replicating a configuration item—whether a program module or subsystem—into multiple instantiations, each of which bears its own unique and appropriate configuration-identification label.

Selecting the appropriate branching model lets the release engineer serve several masters that sometimes have conflicting interests or priorities: the development group, the testing group, and the support group—which represents the product's end users. To determine the adequacy of a branching model, we evaluate its ability to

- maintain a stable base for new development by supporting nightly or continuous integration via builds from the top,
- deliver emergency releases—which contain all necessary fixes and no other changes—to testing and to customers,
- test releases that contain all the necessary fixes and no other changes,
- minimize the impact of emergency releases on new development efforts,
- roll back to a previous production release if necessary,
- support multiple sequential versions in the field, and
- support multiple concurrent versions—such as alternative versions for different platforms or different customers—in the field.

BRANCH-BY-RELEASE MODEL

Conventional wisdom and existing standards tell us to manage a software configuration as a series of successive baselines.[2,3] The *branch-by-release model* of code management instantiates this approach. In this conventional model, the code branches upon a decision to release a new version of the product. The new branch serves as the baseline for continuing development. As Figure 1 shows, the old branch contains the released version—the actual historical baseline reference point. That branch is left behind to wither.

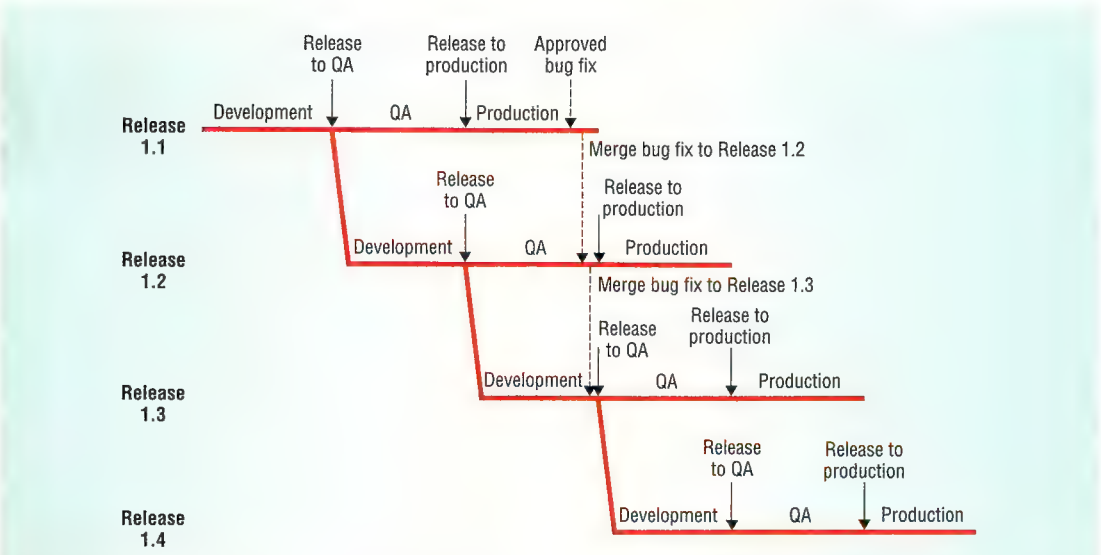

Figure 1. Branch-by-release model. In this conventional model, the code branches when the release engineer delivers a new release. The new branch serves as the baseline for continuing development while the old branch—the code's baseline historical reference—is left behind.

The branch-by-release model appears to provide the series of successive baselines that SCM conventionally requires. It provides a common base for developers to use in making further changes to the code. However, it has two important drawbacks:

- it generally requires serial changes to the code such as sequential check-ins and check-outs, rather than parallel development; and
- it adds complexity and overhead to the support of released versions.

The branch-by-release model is easy to understand—before postrelease bug-fix releases come into play, at least. But making a fix to an earlier release poses many opportunities to lose approved bug fixes or other changes. As long as any version of the product remains supported in the field, the need for a bug fix—and thus an emergency release—remains possible. In such cases, developers must make the fix in that *old* branch and create a new release from it. That's pretty straightforward, but complications arise. Developers must propagate the bug fix down through each subsequent branch to ensure that the bug doesn't reappear in later releases, where it has never been fixed in that release's codeline.

Isolating specific changes and confirming the need to propagate each one to all downstream releases creates an added communication and coordination burden. The environment changes constantly as developers move from one development line to another with each new release. Nor does the model support longer-term development parallel to a release cycle—all code checked out must be checked back in prior to release.

If developers do not check all code back in prior to release, the model does not allow building from the top of the code after release. When the developers check in additional changes after the release, the release engineer must make sure that those untested changes do not find their way into any subsequent new versions, such as emergency releases, created from that line of code. This is particularly problematic because not building from the top drastically increases the risk of incorrect builds for emergency releases. It also undermines the whole baseline concept, because changes never released to the outside have now been introduced in the baseline code.

Finally, the branch-by-release model does not provide a straightforward way to release and maintain multiple concurrent versions in the field.

BUILD-BY-BUG-NUMBER SYNDROME

Using the branch-by-release model leads to the dreaded *build-by-bug-number syndrome.* This occurs when code has been checked in to the old branch after the release, so that the code on the branch no longer matches what was released. The release engineer must handpick the bits of code associated with specific bug fixes that project management or the organization's change control board has decreed necessary for a release. This situation usually arises at the worst time—when the company needs an emergency release. This pressures the team to quickly produce a fix that will remedy an urgent situation.

Tedious at best, and often arduous, the build-by-bug-number process challenges the release engineer to ensure that the release build includes only the pinpointed fixes, all the bits of code needed for each fix, and no other changes. This process often requires the sometimes unwilling participation of the development team, increasing the number of people involved in a tense situation.

How did the project get to this point? Often, the project's management team has determined that changes originally intended for a release could not be included without undue risk. For example, the change might have been seen as potentially destabilizing and there wasn't sufficient time to test it.

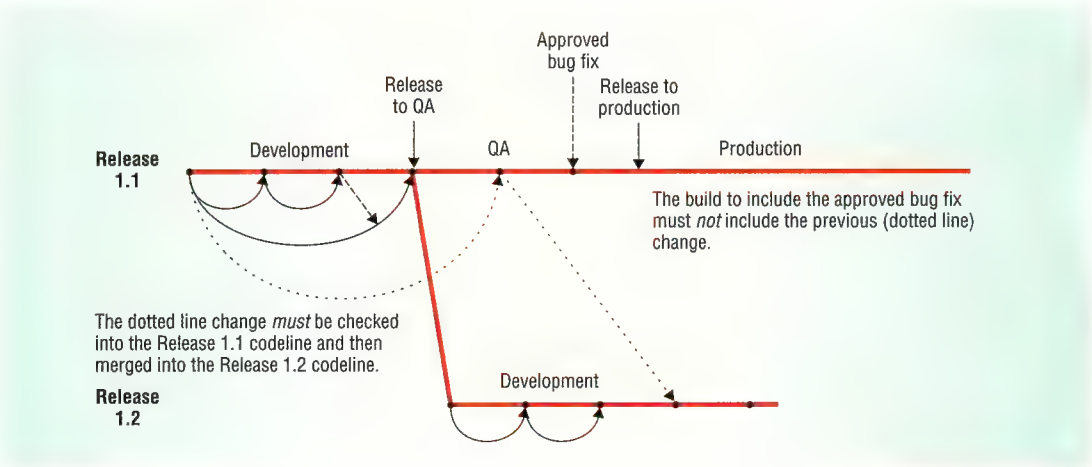

Figure 2. Build-by-bug-number syndrome. This phenomenon occurs when the release engineer must handpick bits of code associated with specific bug fixes, then ensure that only those bits needed for the fixes make it into the release build.

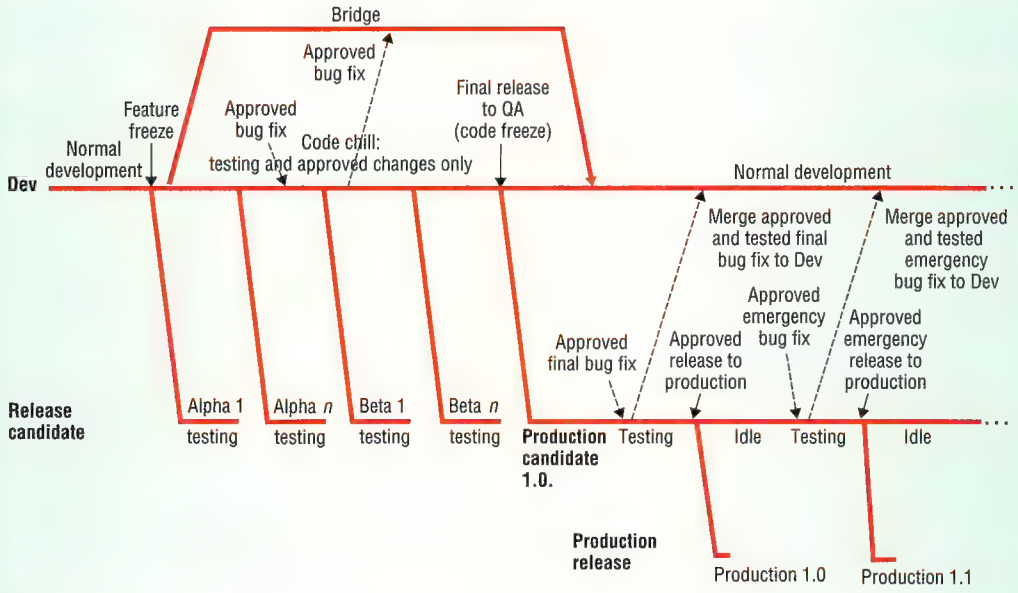

Figure 3. Branch-by-purpose model. The release engineer spins off new branches for specific purposes such as alpha and beta testing, but development work remains on the main development line.

Alternatively, the developers simply could not complete the desired changes in time.

Figure 2 shows the dilemma developers face when confronting this situation. Version control tools require that checked-out code must be checked back in to the same code line from which it was originally taken. If developers do not check the code back in prior to the software product's release, they cannot apply that code directly to the code line being prepared for the next release. The lagging check-in must first be checked into the line it came from, then migrated to the later code line.

Another option requires the developer to abandon those changes and start again on the new development line. Obviously, the developer will oppose this course and the lost work it entails.

As a result, any emergency releases made for that old line must resort to building-by-bug-number. To ensure that only the necessary bug fixes and no other changes—such as lagging check-ins—appear in the emergency release, the release engineer must handpick the specific bits of code associated with the fix. This approach precludes building from the top because doing so would pick up stray changes. Obviously, this scenario invites failure.

In some cases, when poor-quality code requires many postrelease bug fixes, it can lead to abandonment of a disciplined release process in favor of *slipstream* development, as described in the "Slipstream Development" sidebar.

IMPROVED CONFIGURATION MANAGEMENT MODEL

In the *branch-by-purpose model*, shown in Figure 3, release engineering bases the decision to branch on the need to satisfy a specific purpose. Generally, that purpose involves releasing the software and its associated elements, such as documentation, outside the development group. These releases typically mark significant project milestones, such as release to QA for alpha (system) testing, release for beta testing, and so on.

The branch-by-purpose model supports regular releases by design, along with controlled emergency

Slipstream Development

In-house application development groups in information technology organizations often fall into *slipstream* development mode, rather than following a discipline of full product releases. When this happens, the notion of baseline becomes inverted. For one thing, there are no declared product releases. Rather, the *release folders* that these IT groups use may contain just snapshots of the code at particular times, rather than true releases. These snapshots attempt to capture the application's state to establish source code baselines for performing further changes and syncing up the development code line with the actual production code.

They perform this activity *after the fact*, so that what is in the release folders trails the running system. This approach aims merely to capture and archive what's already in production, whatever that happens to be. True SCM, on the other hand, tags and selects in advance the components for building a new version of the application. The release is then built using this information.

This difference is crucial. In the *capture-and-archive model*, the application in production use is the only embodiment of the product's true state and the SCM process follows along, trying to keep track of things and provide developers a reliable baseline to build on. In a true *baseline-and-release model*, the release engineer creates the baseline a priori, using the SCM system to build a new version of the application from previously identified configuration items, and then releases that version, exactly matching the baseline.

Organizations that use a slipstream development-and-delivery process usually rely on the capture-and-archive baselining model. These organizations do not offer true software product releases, per se, just a continuous stream of updates—or patches—to the software in production use. Although this may be convenient for emergency fixes, if overgeneralized it can become the main way every change enters production.

When an organization slips into slipstream development and delivery, it never replaces the entire application. Instead, it updates the application incrementally, which creates some complex problems:

- The continuous-stream model makes it difficult or impossible to recreate past instantiations. There is no way to compare the application as it runs now to the application as it ran several updates back, nor is there a straightforward way to roll back to a previous update stage.
- Testing the evanescent incarnations of the continuous stream can be difficult or impossible—and very expensive. There is rarely a discipline in place to test the full application with each individual changed file—and only that file—and to refuse additional changes until all bugs found in that piece have been resolved and the application has been successfully regression tested with just that changed file.
- The continuous-stream model can also present a challenging security problem. Once a piece of code makes it into production, it won't be replaced until the next fix to that file enters production. If someone slips in a special copy of a piece of code, that piece could stay in production for a long time with no one ever noticing. Although this sort of thing can happen with the replace-the-application model, at least in that case such alterations should leave some tracks in the source code.
- Over time, the always-update, never-replace aspect of the continuous-stream model can lead to the suspicion that some pieces in production weren't built from the source code in the development's current release source folder, such as a live hot fix that didn't make it back into the source folder. This leads to the fear that a full rebuild from the source will not duplicate the running system. If you can't re-create the running system, you're not in control.
- Having reached this uncomfortable position, the development team elects to rebaseline their source code by capturing and archiving the system in production. There is no branching model. This approach simply replaces one line of code with another, at which point the stream starts anew.

The continuous-stream model also implicitly assumes that developers can manage hundreds of files into production well enough, one at a time. Down that path lies madness. Even if the individual files moved along in a completely automated, foolproof process, it is still necessary to manage, test, and track hundreds of interdependent parts as they progress toward production. Even if only 10 percent of the files remain in play between each capture-and-archive baseline event, that's still an impractical number of components to juggle.

releases as required. Further, it avoids the problems caused by branch-by-release as exemplified in the build-by-bug-number syndrome. The branch-by-purpose model will satisfy all typical evaluation criteria.

This model offers the additional benefit of simplifying things for developers. It lets them work in the same environment, the main development branch. This reduces confusion among the development team members about where to make changes, makes for more robust emergency releases, and reduces the team's angst about code chill and code freeze.

We have repeatedly seen reasonably mature and disciplined development teams heatedly resist code freeze in the branch-by-release model because it left them no place to check in their in-progress work. Delaying code chill or code freeze usually compresses the testing cycle, lowering release product quality. Again, forking a new release branch instead of a new ongoing work branch just works better.

On the other hand, managing this model is more complex, primarily because it requires a more sophisticated understanding of SCM and a more sophisticated use of SCM tools. Further, it turns the

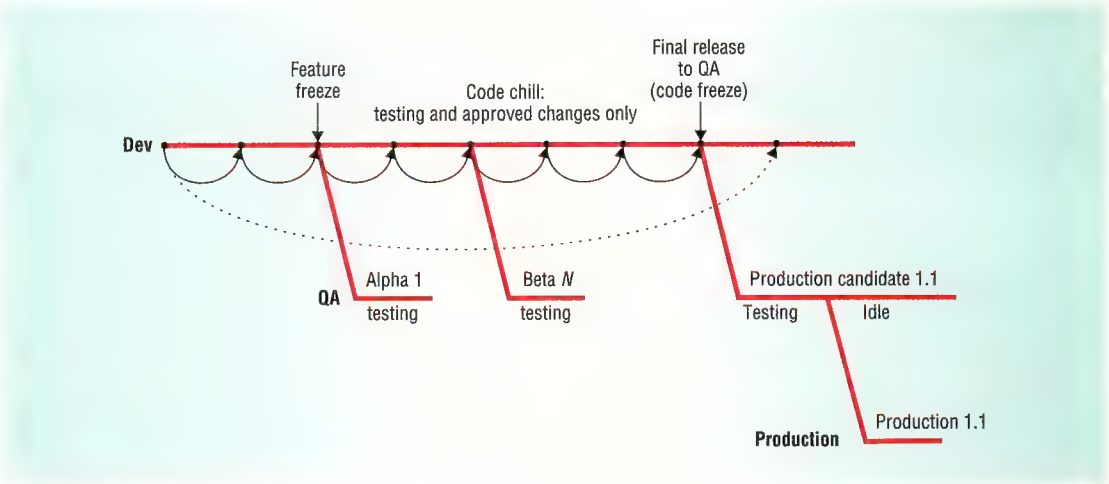

Figure 4. Avoiding the build-by-bug-number tar pit. In the branch-by-purpose model, the developers can just check the code back into the main development line after the release's code freeze, leaving the release code untouched and its integrity intact.

conventional approach on its ear: As before, the release engineer spins off a new branch for the product being released, but development remains on the main dev line.

The product may be released to QA for testing, in which case the release engineer designates a QA branch. Or, once the developers and testers have completed the defect find-and-fix activities, the release engineer can release the product for production use—directly to the user in the case of contracted software, to an operations group in the case of in-house software, or—in the case of shrink-wrapped software—to the manufacturing entity. When releasing the new version, the release engineer creates a production line branch—a special branch whose sole purpose is to support the released version.

The branch-by-purpose model presupposes that the product release cycle includes a *feature freeze* milestone after which developers add no further features and make no more enhancements to existing features without careful change evaluation and control. This milestone marks the entry of development activity into *code chill*. During code chill, developers make fixes in the dev line and periodically send new releases containing the fixes to QA. Each release to QA has it own branch, which makes it possible to verify whether any given bug appeared in an earlier release to QA, and to identify both the release in which a bug is found and the release in which it is fixed.

As the software gains stability through the testing cycle, fewer and fewer fixes become necessary, until code chill gives way to *code freeze*. At the code freeze milestone, the team presumes the product ready for final production use, even though it is only a production candidate that must still undergo final testing to assure its production readiness.

From this point on, the developer applies any fix approved by the change control board (or the project management team) directly to that QA branch and migrates the fix to the main dev line. After testers have verified that the changed code base works properly, release engineering can release the product to production and create a production branch. The team uses this same sort of cycle when a bug found in the field requires an emergency release, as shown in the Production 1.1 build in Figure 3.

SUPPORTING PARALLEL DEVELOPMENT

Today's prevailing climate of rapid development demands parallel development. The software project must proceed with enhancements not intended for the release currently in progress. Even in those rare cases in which rapid development is not de rigueur, the software project may have staff available during code chill and code freeze to implement bug fixes for future releases.

To satisfy either of these needs for parallel development, release engineering can create temporary bridge lines, one for those pending bug fixes and one for the enhancements, or just one bridge line for both bug fixes and enhancements. Then they can merge that line back down into the main dev line after the production candidate releases.

In the branch-by-release model, on the other hand, a problem arises when an in-progress code change must be delayed until after a release. In the case of orphaned check-outs that didn't make it into the release, developers find themselves stuck between a rock and a hard place because they must

- either check the changed code back into the line it came from, which means putting it in the branch with the released product that did not contain it, thereby destroying the integrity of that branch, and then migrate it to the new release branch; or
- check it out all over again from the new branch, redo the changes, and check them back into the new branch.

The branch-by-purpose model obviates this dilemma. As Figure 4 shows, because the main development code line hasn't changed, the developers can just check the code back into the main

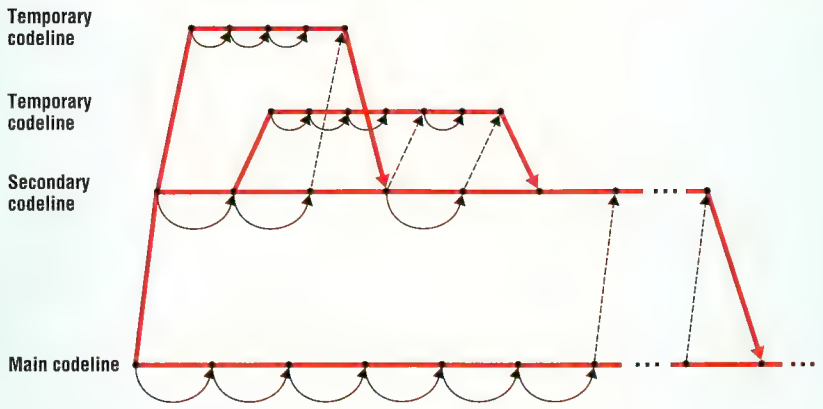

Figure 5. Temporary codeline changes. The release engineer establishes a secondary codeline to handle a major product change. Developers can then branch off one or more temporary codelines from this secondary line and use them to work out minor changes.

line from which it came, and it will be included in a future release's product release cycle. The code base for the released product, Release 1.1, remains untouched because it resides on its own branch.

Multiple groups and subprojects

Branch-by-purpose more readily accommodates the need to have multiple groups working on multiple subprojects in parallel. Branching can establish a fully replicated environment in which developers can modify code and test changes without impeding other developments occurring simultaneously in the code base. Creating small, purposeful, alternate development lines lets developers check their code into the appropriate temporary line without affecting the main line.

The parallel development concept recognizes the need to frequently merge changes from one line into another. Because the branch-by-purpose model anticipates that need, the SCM and release engineering processes can manage merges by design, making them as painless as possible.

In contrast, the branch-by-release model operates from the underlying assumption that releases are linear and sequential, each subsequent release flowing immediately from its predecessor. Thus, changes made to an existing release of the current baseline—or its predecessor, if still in the field—are somewhat unnatural, and making these changes requires special effort to ensure that

- any re-release of an old release contains just the desired changes—typically, bug fixes—and no others, and
- the release engineer properly merges these changes into the new, in-progress release and into all releases between the old release and the current release.

In the branch-by-purpose model, on the other hand, developers work from the underlying assumption that they are creating a single main line of code and other artifacts and that, at predetermined times, release engineering will create a separate branch when it releases a new version of the product. This approach enables frequent builds from the top on the main line of code that thus include all available changes. The best practices of today's software companies emphasize frequent integrations via builds from the top. Such builds integrate changes bit by bit, as testers and developers check them in. Builds can be nightly or continuous. This practice assures continual integration and avoids the big bang that lots of colliding changes can cause when the project team delays integration until the release cycle's end.

Having a fully replicated environment for modifying code enables continuous or nightly builds of each line and further testing of those changes, without impeding the main line's development. It also lets multiple developers make multiple changes to the code for multiple purposes.

While project teams can perform branching for parallel development at very fine granularity levels—down to each bug fix—they use two common modes to create branches for alternate development lines:

- when one or several developers will be working on a small-to-middle-sized change or related group of changes over several weeks, or
- when many developers must work on a large change over several months.

Variations of the branch-by-purpose model can accommodate each mode.

Codeline changes

When testing shows that changes work as expected, the developers merge the changes into the main development line. As Figure 5 shows, developers also periodically merge main codeline bug fixes and other changes to the temporary line if it must be maintained for more than a few weeks.

Whether or not developers periodically merge the main line's changes, they must merge the changes up to the temporary line before merging it

back down to the main line. Integration takes place on the temporary line, assuring that its changes are compatible before inserting them into the main line, which the rest of development depends on.

To make larger, more complex changes, developers establish a secondary codeline to accommodate a major product change, such as a rearchitecting effort. They merge the changes to the main codeline up to their secondary line every week or so.

Figure 5 shows two temporary lines branching off the secondary line to let developers work out the kinks in some small- or medium-sized changes that depend on changes in the secondary line.

Despite its wide use and its appearance of mirroring the conventional standard of sequential baselines, the branch-by-release model imposes unnecessary burdens on the software project team in supporting released versions of software while developing new releases. Error-prone, it also often requires manual handling of fixes to released software. Further, its fundamental flaws include unnecessary complexity in managing post-release code fixes and unnecessarily orphaning changes in progress, which cannot be successfully injected prior to a scheduled release date.

The branch-by-purpose model, on the other hand, avoids these pitfalls. In addition, it provides a structured mechanism for managing multiple lines of code, supporting the organization's need to deliver multiple releases as well as development's need to shorten cycle time by using parallel development and continuous integration. ■

Acknowledgments

We thank Matt Foley and Adonica Gieger for their helpful comments in reviewing this article.

References

1. P.H. Feiler, *Configuration Management Models in Commercial Environments*, tech. report CMU/SEI-91-TR, Software Eng. Inst., Carnegie Mellon University, Pittsburgh, 1991.
2. *ANSI/IEEE Std. 828-1983, IEEE Standard for Software Configuration Management Plan*, IEEE Press, Piscataway, N.J., 1983.
3. *ANSI/IEEE Std. 1042-1987, IEEE Guide to Software Configuration Management Content*, IEEE Press, Piscataway, N.J., 1987.
4. R. Conradi and B. Westfechtel, "Version Models for Software Configuration Management," *ACM Computing Surveys*, vol. 30, no. 2, 1998, p. 233.
5. I. Sommerville, *Software Engineering*, Addison-Wesley, Reading, Mass., 1996, p. 657.

Chuck Walrad is managing director of Davenport Consulting. Her research interests include strategies for the propagation and practical implementation of best practices in software engineering. Walrad received an MS in theoretical linguistics from the University of California, San Diego. She is a member of the IEEE and the ACM. Contact her at cwalrad@daven.com.

Darrel Strom is a senior software engineer at Expert Support. His research interests include developing software factories that automate the critical processes required to build and deliver software products reliably. Strom received a BS in computer science from the University of Southern Mississippi. He is a member of the IEEE and the ACM. Contact him at dstrom@xs.com.

PICO: Automatically Designing Custom Computers

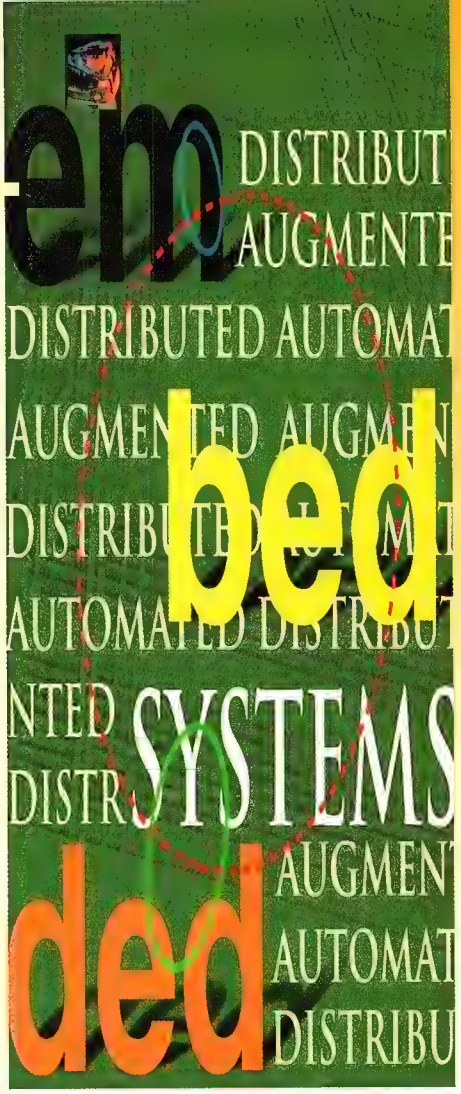

The PICO project automates the design of optimized, application-specific embedded computer systems to meet the demands of innovative smart products that require varying combinations of performance and cost.

Vinod Kathail
Shail Aditya
Robert Schreiber
B. Ramakrishna Rau
Darren C. Cronquist
Mukund Sivaraman
Hewlett-Packard Laboratories

Embedded computers are everywhere—in video games, DVD players, TV sets, printers, scanners, cell phones, cars, and now even in smart robotic vacuum cleaners, lawnmowers, and virtual pets. Computers have displaced many analog circuits in photography, video, and telephony. The advent of system-level integration (SLI) foreshadows an era of yet more growth in the number and variety of innovative smart products and their embedded computers.

Such smart products demand varying combinations of performance, cost, and power. When a product mandates high performance, often the challenge is to lower cost to a level the market will accept. Whereas specialization increases performance and reduces cost, customization permits specialization when no adequately specialized, off-the-shelf design is available. Automation of embedded computer design would enable customization by lowering the barriers of design time, designer availability, and design cost, thereby unleashing the predicted explosion in smart products.

PICO ARCHITECTURE SYNTHESIS SYSTEM

The PICO (program in, chip out) project at HP Labs automates the design of optimized, application-specific computer systems. PICO uses an application written in C to architect a set of high-quality system designs that trade cost for performance.

As Figure 1 shows, a PICO system design contains one EPIC/ VLIW (explicitly parallel instruction computing/very long instruction word) processor[1] and an optional nonprogrammable accelerator (NPA) subsystem consisting of one or more NPAs, both connected to a two-level cache subsystem that, in turn, connects to the system bus. Each NPA is customized to execute a compute-intensive loop nest that would otherwise have been executed on the VLIW.

PICO generates the most cost-effective combinations of these subsystems to provide several high-quality system designs at varying points on the cost-performance tradeoff curve. PICO emits structural Verilog/VHDL for the hardware components, modifies application code to include software interfaces to the generated hardware, and retargets the compiler, assembler, and simulator to the custom VLIW processor.

Skeptics often assume that automated design must emulate human designers who can invent new solutions to problems. PICO's approach, however, is to automatically pick the most suitable designs from a large, well-engineered space of designs. In practice, it would be unrealistic for all designs to be

0018-9162/02/$17.00 © 2002 IEEE

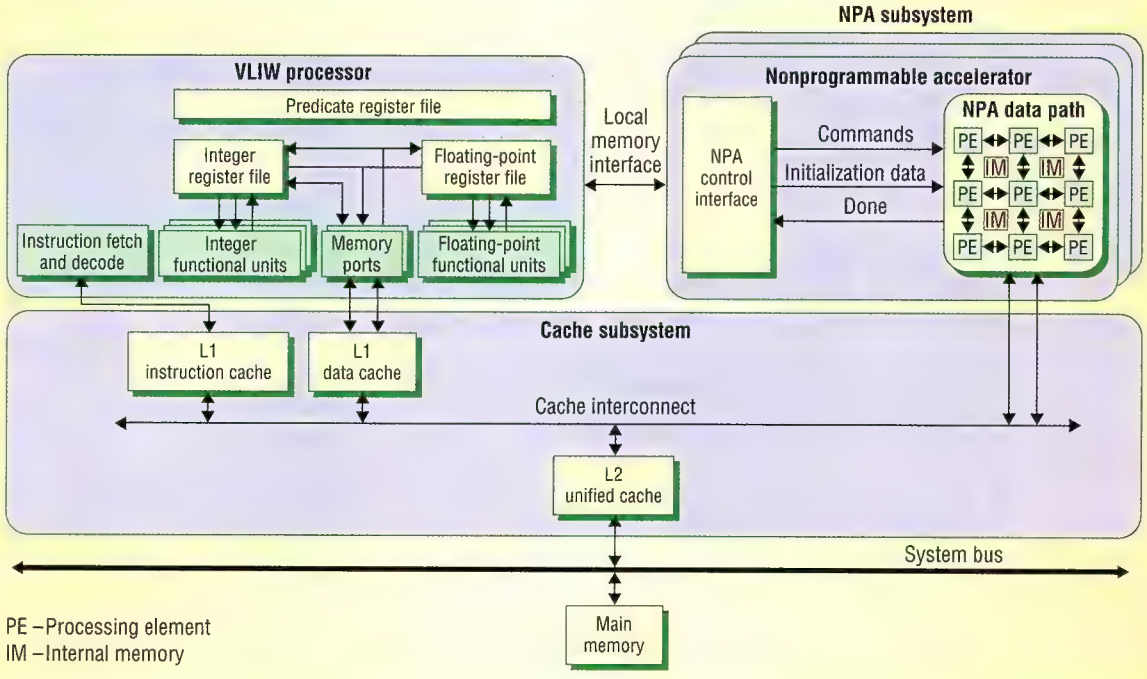

Figure 1. PICO's system-level architecture template. A system consists of three subsystems: a custom EPIC/VLIW processor; a custom, two-level cache hierarchy; and, optionally, an NPA subsystem that includes one or more custom NPAs.

preconstructed, so PICO's framework-based, hierarchical design methodology creates designs on demand during design space exploration.

FRAMEWORK-BASED AUTOMATION

Figure 2 shows PICO's framework-based design automation strategy. A framework consists of

- a parameterized architectural template that defines the space of designs to be considered,
- a spacewalker with a strategy for exploring the design space,
- a constructor that can construct every design in this space using components from a component library, and
- an evaluator that can measure the quality of any such design.

These elements together provide the basis for automatically identifying an approximate Pareto set—a set of designs, each of which is better than any other design in at least one measure of quality.

Template. The template defines parameters representing the design space and a set of rules and constraints that must be honored. Within the template, some aspects of the design, such as the presence of certain modules and how they connect, are predetermined. Architectural parameters specify other aspects of the design. For example, the number of memory ports in a processor might be a parameter in a processor design space. Once the parameters have been specified, a detailed construction of the design determines many of its other attributes. For example, the VLIW constructor determines a VLIW processor's instruction format.

When we have established an optimal or near-optimal algorithmic way of determining the design details, we view their definition as part of the construction task. When we have no clear way of determining a design's important aspects and must use a heuristic search to determine good values, we view these attributes as parameters. The number of parameters should be relatively small to be manageable. A specification is the set of values to which the parameters are bound and corresponds to a unique design.

Spacewalker. The set of all allowed parameter value combinations in the template defines the specification space. The spacewalker explores this space, looking for the Pareto-optimal designs. The template may fix the range of allowed values for a parameter, the user may set the range, or the spacewalker may determine the range through a preliminary examination of the application.

Using the quality metrics the evaluator provides, the spacewalker determines whether any previously examined design *eclipses* the new design—the new design is either equal or inferior to the previous design in all respects. If not, the new design is added to the Pareto set, discarding all previous designs in the set that the new design eclipses.

If a design space is too large to search exhaustively, we use a manifold strategy to search the design space more intelligently. The spacewalker uses heuristics to examine designs that are likely to be Pareto-optimal while avoiding the exploration of uninteresting designs. The goal is to find most of the Pareto-optimal designs while having examined only a small fraction of the specification space.

The spacewalker also maintains a repository of

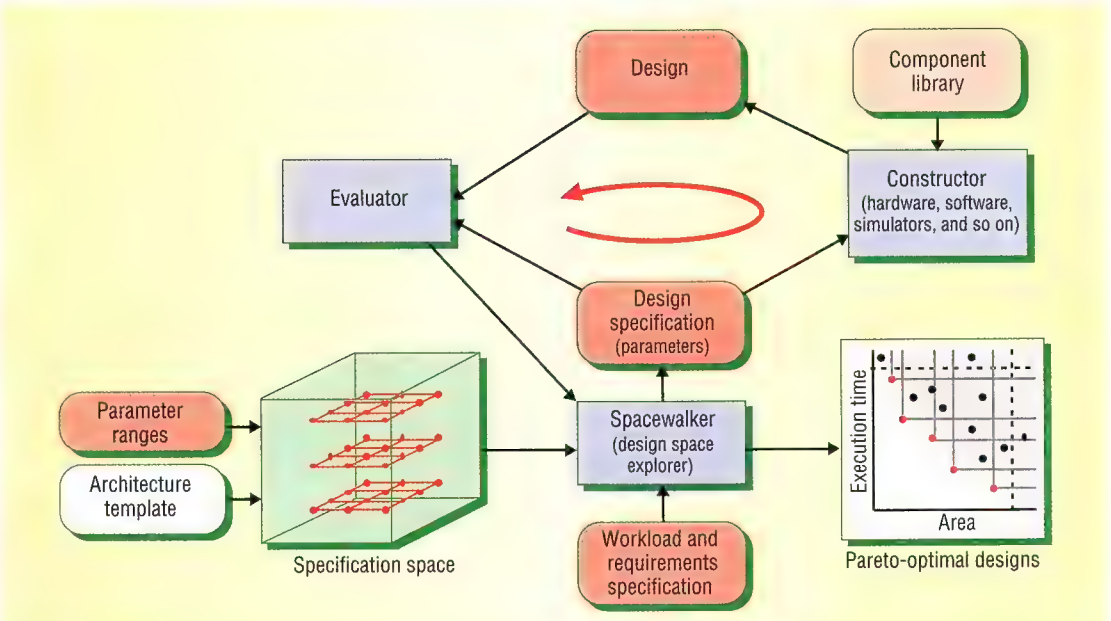

Figure 2. PICO's framework-based automation strategy. The PICO framework consists of a template, spacewalker, constructor, evaluator, and component library.

previously explored designs to eliminate redundant examination of the same design. Analytical cost and performance modeling can help avoid construction and evaluation costs entirely. Finally, we use the basic divide-and-conquer paradigm, which in this context takes the form of hierarchical design and hierarchical design space exploration.

Constructor. Design construction derives a detailed design from the specification provided by the spacewalker. The spacewalker specification is, of necessity, an abstract description of the machine to be synthesized. The constructor fills in myriad details using the abstract description to create the best possible—and usually least costly—machine.

To construct the design, the constructor automatically assembles lower-level components chosen from a component library. Sometimes, these components are parameterized with respect to their structural properties such as bit width; once the parameters are specified, a generator automatically instantiates the specific component. The components themselves are designed, optimized, and characterized manually for properties such as area, power, and latency.

Evaluator. The framework also requires an evaluator that computes various metrics to assess the detailed design, the design specification, or both. In PICO, the cost of a design is measured in terms of gates or silicon area, and its performance is measured in terms of the number of cycles to execute the application via a combination of static estimation and simulation.

HIERARCHICAL DESIGN

To limit design complexity, designers decompose systems into subsystems that interact in a limited way. Identifying the best subsystem designs and then combining them to form complete systems greatly reduces the system design space. A hierarchical design framework formalizes this intuitive approach.

A hierarchical system design framework decomposes the system into several subsystem frameworks. Each framework must include a template that defines the subsystem specification space, a spacewalker, a constructor, and an evaluator. Developers must define a system framework's specific decomposition strategy while designing the system framework, prior to automating its design. System composition, on the other hand, occurs during design space exploration.

System decomposition involves not only the system template's structural decomposition, but also decomposition of the constructor and the evaluator. As good software engineering practice demands, the subsystem constructors are typically already in place.

However, the challenge is decomposing the system evaluator. Subsystem evaluators can use different metrics than the system evaluator, but they must be good indicators of system-level performance. Specifically, the evaluators should reject poor subsystems without omitting subsystems that are system-level constituents of Pareto designs.

The key point in hierarchical design is that the design space at each level in the hierarchy is restricted to systems that can be built only from the Pareto-optimal subsystems. This restriction greatly reduces the size of the system design space that must be explored.

Spacewalking without decomposition

Consider, for example, a computer system template that consists of a processor and a data cache. Based on the specified parameters and ranges for these elements, suppose the system specification space holds 5,000 designs. An exhaustive spacewalk examines all 5,000 system specifications to find, for

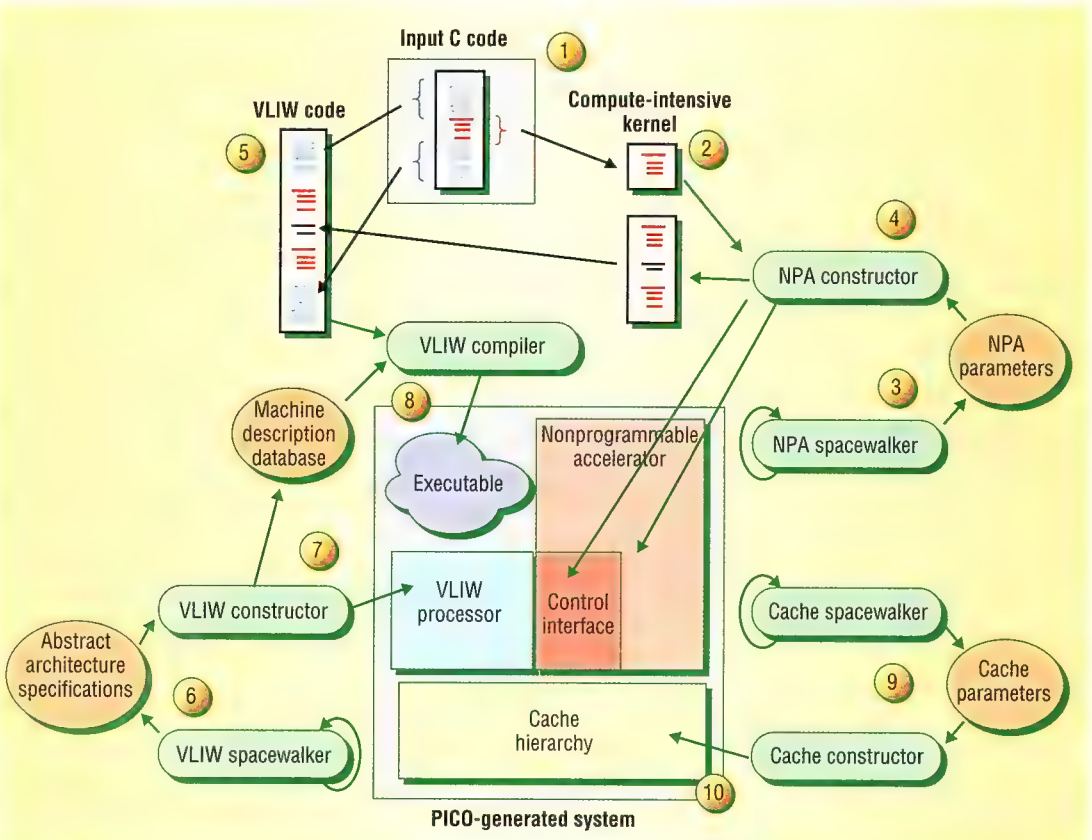

Figure 3. PICO's hierarchical design flow. The NPA, VLIW, and cache subsystem frameworks use the input C application to produce Pareto-optimal subsystems designs that are composed to produce Pareto-optimal system designs.

example, 15 Pareto designs. These specifications comprise only 0.3 percent of the evaluated systems. Hierarchical design aims to reduce the number of designs examined without missing any Pareto points.

Spacewalking with decomposition

Consider a decomposition of the above system framework into a hierarchy consisting of a processor and a cache framework, each with its own set of parameters. Some system parameters become parameters of only one subsystem; others affect more than one subsystem. The exploration ranges of the subsystem parameters are derived from those of the system parameters. In addition, decomposition liberates some previously dependent system attributes, such as cache ports, and makes them independent subsystem parameters.

In this example, suppose there are 100 processor and 100 cache designs in their respective specification spaces, each containing 10 Pareto-optimal designs. Then, an exhaustive spacewalk examines only 300 designs—100 processor, 100 cache, and 100 system designs—to find the 15 Pareto designs, raising the exploration efficiency to 5 percent.

System composition and evaluation

In a well-formed system, various subsystems must obey certain interface constraints. In this example, one of the interface constraints is that processor load-and-store units—a processor subsystem parameter—must be equal to the cache-ports parameter of the cache subsystem. Such interface constraints prevent subsystem designs from being interchangeable: A one-port data cache cannot be used in place of a two-port data cache, for example. Spacewalking for subsystems must account for this, or it will too aggressively remove subsystem designs that are necessary for Pareto-optimal system designs.

The system-level constructor must ensure that it builds and evaluates only valid combinations of subsystems. One possibility is to perform *validity filtering*, forming all possible combinations of Pareto subsystems and testing whether they satisfy the interface constraints to weed out invalid subsystem design combinations.[2]

When an interface constraint is a simple equality or inequality involving subsystem parameters, PICO avoids creating invalid systems altogether, using Pareto sets *indexed* by parameters to construct the valid systems directly.

PICO DESIGN FLOW

Figure 3 shows PICO's hierarchical design space exploration. After inputting a C application containing one or more compute-intensive loop nests or kernels (1), PICO first identifies and extracts each kernel (2). Then the PICO-NPA spacewalker repeatedly specifies an NPA, retaining only the best NPA designs (3). For each specification, the NPA constructor transforms the kernel to the requisite level of parallelism and main memory bandwidth

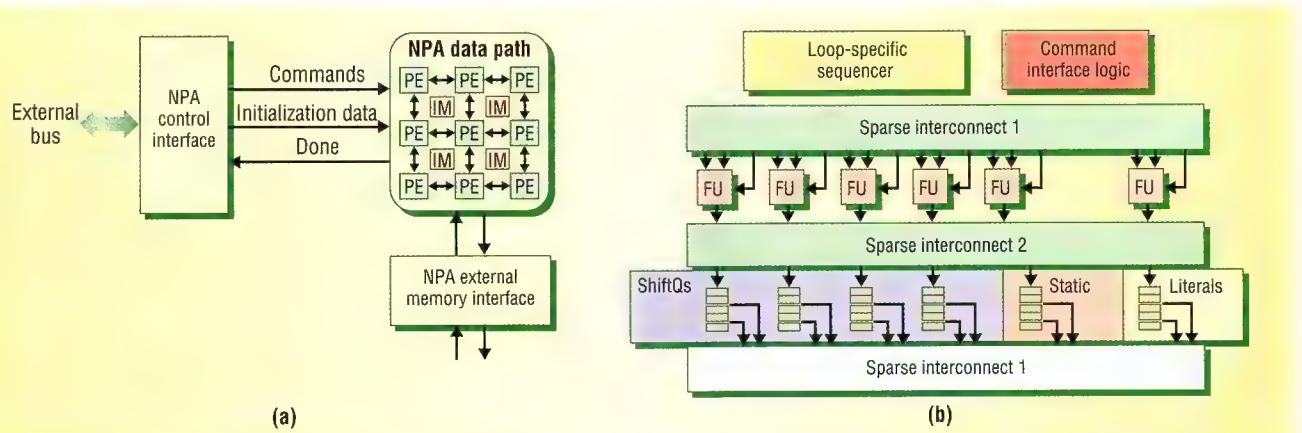

Figure 4. The NPA template. (a) An NPA consists of a one- or two-dimensional array of customized processing elements (PEs), with only synchronous parallel nearest-neighbor communication, and internal memories (IMs). (b) The given loop nest computation largely determines each PE's structure, which consists of an interconnected network of functional units and distributed register structures (ShiftQs).

and generates the register transfer level (RTL) design for the NPA along with the VLIW code that will repeatedly initialize and invoke the NPA (4).

At this point, using the combined VLIW code (5), the PICO-VLIW spacewalker repeatedly specifies a VLIW processor and retains only the best VLIW designs (6). For each specification, the VLIW constructor designs the VLIW processor's architecture and microarchitecture and emits an RTL design (7). In addition, the constructor generates a machine description of this processor for the Elcor VLIW compiler, which compiles the modified application to the VLIW processor (8). The cache spacewalker repeatedly specifies and evaluates cache subsystem configurations, retaining only the Pareto-optimal ones (9).

Finally, the system-level PICO spacewalker (not shown in the figure) combines compatible VLIW, cache, and NPA designs into Pareto-optimal system designs (10). During this process, and for each cost or performance level, PICO performs hardware-software partitioning and codesign by determining which kernels should be implemented as NPAs rather than as software on the VLIW.

NPA DESIGN

PICO-NPA accepts a loop nest in C, along with a range of performance requirements and available external memory bandwidth, and produces a Pareto set of NPAs customized for the given loop nest.[3] For each Pareto design, PICO emits structural Verilog/VHDL that defines the NPA at the register transfer level, as well as requisite initialization and invocation code for the external host processor to execute.

Template

The NPA template consists of an array of customized processing elements (PEs) with only synchronous parallel nearest-neighbor communication, as Figure 4a shows. Further, the design can contain internal memories (IMs), an interface to external memory, and a memory-mapped host processor interface. As Figure 4b shows, each PE has a loop-specific instruction sequencer and a data path consisting of

- functional units (FUs) with customized widths;
- distributed register structures with individual read and write access from and to any register, as opposed to addressable register files; and
- sparse interconnects customized to the loop nest.

The data path lacks centralized instruction storage, distributing the loop nest's operations across the FUs instead.

Constructor

First, a suite of loop-nest transformations and optimizations determines the placement of data in external memory, internal memory, or registers. Then the constructor tiles the iteration space to minimize the number of registers required while maximizing available external memory bandwidth utilization. The constructor schedules the loop iterations in time and maps them to PEs using the available parallelism to meet the given performance specification. Additional low-level optimizations follow, both standard (common subexpression elimination, dead code elimination, strength reduction) and novel (data-width inference and clustering).

Then, the constructor allocates a least-cost set of FUs capable of executing the operations in the optimized loop nest at the desired performance level given by its initiation interval (II) by solving an integer linear program. The constructor performs modulo scheduling to determine each operation's issue time and to bind each operation to an FU, so that each operation executes every II cycles without conflicting for resources.

The constructor uses various heuristics that cluster similar width operations and maximally share storage and interconnect, minimizing hardware costs prior to and during the scheduling and binding phase. ShiftQ[4]—a novel hardware structure consisting of registers and switches—buffers and transports operands between FUs. A dedicated ShiftQ

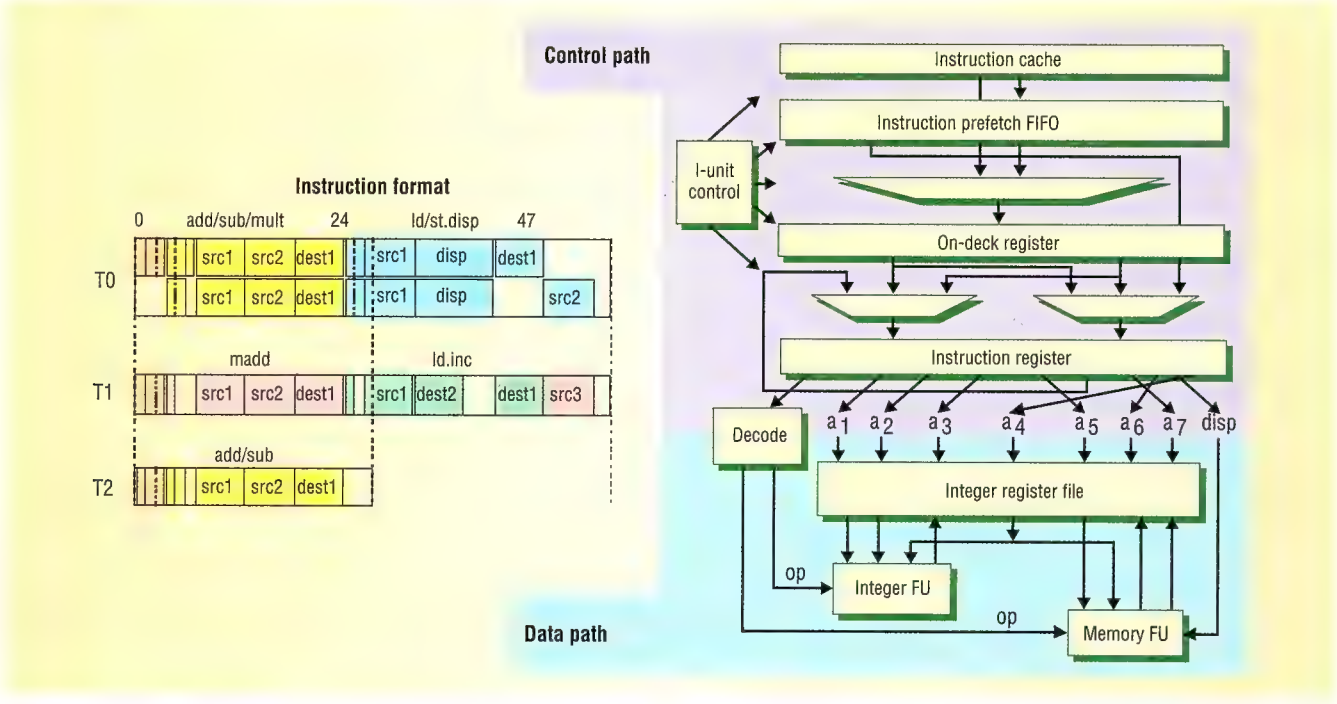

Figure 5. VLIW processor template. Although it fixes certain aspects of the design, the design allows flexibility elsewhere. For example, every VLIW processor contains an integer cluster consisting of an integer register file and a set of integer functional units connected to it. Optionally, the processor can contain a floating-point cluster with the same general structure.

for each FU minimizes the storage required for the results produced by the FU. The constructor creates connections from ShiftQ registers to FU inputs as needed, resulting in a sparse interconnect structure.

The constructor generates the instruction sequencer and creates multiple copies of the PE, interconnecting them in the geometry the spacewalker specifies. The constructor also generates internal memories, interfaces to the host processor and external memory, and provides arbitration and stalling circuitry, if needed. Finally, the constructor emits structural Verilog/VHDL for the NPA hardware and generates the C code that repeatedly invokes the NPA hardware after making appropriate initializations. This code is compiled onto the host processor along with the remainder of the application.

The NPA construction process permits simulation at several intermediate steps. By comparing these simulation results with those of the original source code, developers can detect errors in the input specification as well as errors that the PICO software introduces. In addition to the RTL artifact, the constructor also generates an RTL testbench and memory simulation models for block-level RTL verification and a cycle-accurate C model to support system-level verification.

Evaluator

The NPA cost evaluation uses the parameterized formulas for area and gate count attached to each component in the macrocell library to estimate the design's chip area and gate count. These formulas have been calibrated for a specific ASIC design-flow and process technology. The constructor designs the NPA to a steady-state performance supplied as a parameter. To compute the overall NPA performance, the evaluator makes adjustments for the time required to fill and drain the pipeline of PEs and for the delay anticipated due to memory-induced stalls.

VLIW PROCESSOR AND COMPILER DESIGN

PICO-VLIW uses a C application, with test data, to automatically produce a set of Pareto-optimal, custom VLIW processor designs. For each design, PICO-VLIW architects the processor—including the instruction format, the execution data paths, and the instruction unit—and emits structural Verilog/VHDL for it. In addition, it retargets the Elcor compiler for VLIW architectures to the newly designed processor.[5]

Template

The PICO-VLIW template shown in Figure 5 encompasses a broad class of EPIC/VLIW processors having advanced architectural and microarchitectural features.[1] Their operation repertoire is a customized subset of the HPL-PD operation set,[6] optionally augmented by user-defined operations. The processors can issue multiple operations per cycle executing on multiple FUs. In addition to the global address space that the main memory and the

cache subsystem represent, a processor can use distinct load and store operations to access one or more local memories.

HPL-PD specifies four files for organizing the registers: integer, floating-point, predicate, and branch. FUs that do not need simultaneous access can share register file ports. To reduce code size, two or more FUs that cannot issue operations simultaneously can share the same bit positions in the instruction format. This format can consist of several instruction templates of different lengths to reduce the number of explicit no-op operations.

Spacewalker

The VLIW spacewalker emits a high-level architecture specification that specifies the types and sizes of register files, the operation repertoire as a collection of operation groups, and an abstract specification of the processor's instruction-level parallelism as mutual exclusions between operation groups. Operations in two different operation groups execute in parallel unless there is a mutual exclusion between the two groups. Furthermore, operations within an operation group cannot be issued in parallel. In the current implementation, the operation groups are instances of four operation types—the integer, floating-point, load-and-store, and branch operations—and there are no explicit mutual exclusions between operation groups.

The VLIW spacewalker customizes the operation repertoire to the application: Any operation in HPL-PD repertoire that the application does not use can be deleted. Alternatively, we can retain a core set of operations if required for greater generality.

The VLIW spacewalker's specification space can be quite large, and the evaluation of each design point can be time-consuming because it involves recompiling the full application program. Thus, exhaustive search is impractical, and the spacewalker adaptively limits its specification space search based on information gleaned from the points it has already examined.[7] The key idea is to make incremental changes to the current design's parameters and to terminate the current search direction if all the current design's neighbors are non-Pareto designs.

Constructor

The VLIW constructor first generates the least expensive data path it can, consistent with the concurrency level the abstract architecture specification demands. The constructor uses the specified mutual exclusions to minimize the number of logical FUs required and to maximize register port sharing between them. The constructor may use a single physical FU to implement the union of operations in mutually exclusive operation groups. Or, it may use multiple physical FUs that share register file ports.

The constructor then designs an instruction format for the processor using mutual exclusions and compiler feedback. The instruction format consists of a number of variable-length instruction templates that judiciously balance the total code size against the complexity of the instruction decode and distribution network.

The spacewalker adaptively limits its specification space search based on information gleaned from points it has already examined.

The constructor uses mutual exclusions to design a basic instruction format that supports the requisite level of concurrency. It also generates an Elcor machine description, compiles and schedules the application, and gathers statistics on which sets of operations are frequently issued together. It uses these statistics to augment the basic instruction format with additional templates for operations that are frequently issued together in order to reduce the total code size.

The constructor next generates an instruction unit consisting of the instruction prefetch, alignment, and decode hardware customized for this instruction format using the schema specified in the VLIW template.

At this point, the detailed design is complete, and the constructor generates structural Verilog/VHDL. It also creates the final machine description needed to retarget the compiler, assembler, and simulator.

Evaluator

The VLIW evaluator estimates chip area and gate count in the same way as the NPA evaluator. To estimate the application's processor runtime, the evaluator multiplies each basic block's schedule length by its profiled execution frequency and sums the total over all basic blocks. This method produces an accurate performance estimate for a statically scheduled processor ignoring stall cycles due to cache misses. The cache subsystem evaluator is responsible for computing the additional stall cycles due to cache misses.

The evaluator also reports the generated object code's size to evaluate its incremental effect on the cost of main memory, which is charged to the VLIW cost. The code size also affects the evaluation of instruction cache (I-cache) and unified cache (U-cache) misses in the I-cache framework.

SYSTEM DESIGN

System framework decomposition is valid only if the various subsystem frameworks can be inde-

Framework-based automation offers a powerful methodology for automating the design of complex processors and computer systems.

pendently constructed, evaluated and then combined to give a reasonable approximation to the overall system Pareto. PICO's system-level decomposition strategy, therefore, hinges upon two important factors.

First, subsystem compatibility is captured by using a set of additional parameters for each subsystem that participates in interface constraints during system composition. Given a family of parameterized subsystem Pareto sets, the system-level PICO spacewalker combines only compatible VLIW, cache, and NPA designs into a system-level Pareto set. In the process, PICO determines which kernels should be implemented as NPAs rather than as software routines on the VLIW.

Second, PICO assumes that each subsystem can be independently evaluated for cost and performance using its parameters. This assumption is valid under certain design constraints that PICO enforces, and it has been verified both analytically and experimentally. For example, the evaluation of the memory hierarchy is broken into separate evaluations for the data cache (D-cache), I-cache, and U-cache. This decomposition is valid if the U-cache includes all data contained in the I-cache and D-cache.

PICO uses dilation, an empirical parameter that is the ratio of the compiled application code size on the given processor with respect to a fixed reference processor, to capture the effect of varying the instruction format of various VLIW processors over the I-cache and the U-cache instruction misses.[8] Processors with various dilations are evaluated independently and then matched with cache subsystems with the same dilation without the need to evaluate each processor-cache pair separately.

Framework-based automation offers a powerful methodology for automating the design of complex, high-level structures such as processors and computer systems. It has been crucial to PICO's success in designing NPAs and VLIW processors. We strongly suspect that automated design is possible only with this sort of methodology.

The restrictions that a parametric template places on the design space remove enough variability to make automation possible; the existence of a template in turn makes designing automatic spacewalkers, constructors, and evaluators possible. The nature and number of parameters in the template determine the framework generality and flexibility. Frameworks can be quite varied. The template for one framework might contain RISC processors and DSPs, while another might contain a vector processor.

Hierarchical design methodology, on the other hand, makes the problem tractable. Our experience with hierarchical design indicates that decomposing the evaluator presents a very challenging problem. Developers can easily decompose a template to identify the interface constraints and parameters, and to design spacewalkers and constructors for the resulting subsystems. Engineering acceptably accurate subsystem evaluators requires judicious decomposition. We believe that this is the most important research problem in hierarchical design. Until this problem is solved in an adequately general manner, developers of complex systems might need to restrict themselves to frameworks that easily lend themselves to hierarchical evaluation.

Engineering disciplines tend to go through fairly predictable phases: ad hoc, formal and rigorous, and automation. When the discipline is in its infancy and designers do not yet fully understand its potential problems and solutions, a rich diversity of poorly understood design techniques tends to flourish. As understanding grows, designers sacrifice the flexibility of wild and woolly design for more stylized and restrictive methodologies that have underpinnings in formalism and rigorous theory. Once the formalism and theory mature, the designers can automate the design process. This life cycle has played itself out in disciplines as diverse as PC board and chip layout and routing, machine language parsing, and logic synthesis.

We believe that the computer architecture discipline is ready to enter the automation phase. Although the gratification of inventing brave new architectures will always tempt us, for the most part the focus will shift to the automatic and speedy design of highly customized computer systems using well-understood architecture and compiler technologies. ■

Acknowledgments

We thank past members of our group: Scott Mahlke, Mike Schlansker, Santosh Abraham, Greg Snider, Sadun Anik, and Richard Johnson. Alain Darte and Lothar Thiele made certain key theoretical contributions. The following research interns accelerated PICO software development: Frederic Vivian, Bruce Childers, Marnix Arnold, Matthew Jennings, Roderic Rabbah, Timothy Sherwood, Rajiv Ravindran, Guillaume Huard, Luca Ceresoli, Kevin Fan, and Benoit Meister. We thank Henk Corporaal for providing the code for the attractive GUI to his MOVE infrastructure, from which we

developed PICO's GUI, and Wen-mei Hwu and his IMPACT group for providing Elcor's machine-independent front end.

References

1. M.S. Schlansker and B.R. Rau, "EPIC: Explicitly Parallel Instruction Computing," *Computer*, Feb. 2000, pp. 37-45.
2. S.G. Abraham and B.R. Rau, "Efficient Design Space Exploration in PICO," *Proc. Int'l Conf. Compilers, Architecture, and Synthesis for Embedded Systems* (CASES 2000), ACM Press, New York, 2000, pp. 71-79.
3. R. Schreiber et al., "PICO-NPA: High-Level Synthesis of Nonprogrammable Hardware Accelerators," *J. VLSI Signal Processing*, vol. 31, 2002, pp. 127-142.
4. S. Aditya and M.S. Schlansker, "ShiftQ: A Buffered Interconnect for Custom Loop Accelerators," *Proc. Int'l Conf. Compilers, Architecture, and Synthesis for Embedded Systems* (CASES 2001), ACM Press, New York, 2001; http://hplabs.hp.com/research/itc/car/Templates/carpapers-hpl.html#2001.
5. S. Aditya, B.R. Rau, and V. Kathail, "Automatic Architectural Synthesis of VLIW and EPIC Processors," *Proc. Int'l Symp. System Synthesis* (ISSS 99), IEEE CS Press, Los Alamitos, Calif., 1999, pp. 107-113.
6. V. Kathail, M.S. Schlansker, and B.R. Rau, *HPL-PD Architecture Specification: Version 1.0*, tech. report HPL-93-80R1, Hewlett-Packard Laboratories, Palo Alto, Calif., 2000; http://www.hpl.hp.com/techreports/93/HPL-93-80R1.html.
7. G. Snider, *Spacewalker: Automated Design Space Exploration for Embedded Computer Systems*, tech. report HPL-2001-220, Hewlett-Packard Laboratories, Palo Alto, Calif., 2001; http://www.hpl.hp.com/techreports/2001/HPL-2001-220.html.
8. S.G. Abraham and S.A. Mahlke, "Automatic and Efficient Evaluation of Memory Hierarchies for Embedded Systems," *Proc. 32nd Ann. Int'l Symp. Microarchitecture* (MICRO 99), IEEE CS Press, Los Alamitos, Calif., 1999, pp. 114-125.

***Vinod Kathail** is a principal research scientist and manager in the Compiler and Architecture Research Program at Hewlett-Packard Laboratories. His research interests include parallel computer architectures, compilers, programming languages, and automated design of custom computer systems. He received an ScD in electrical engineering and computer science from MIT. He is a member of the IEEE and the ACM. Contact him at vinod_kathail@hp.com.*

***Shail Aditya** is a senior research scientist in the Compiler and Architecture Research Program at Hewlett-Packard Laboratories. His research interests include programming language design, multithreaded and VLIW compilers and processor architectures, embedded system design, and design automation. He received a PhD in electrical engineering and computer science from MIT. He is a member of the IEEE Computer Society. Contact him at aditya@hpl.hp.com.*

***Robert Schreiber** is a principal scientist in the Compiler and Architecture Research Program at Hewlett-Packard Laboratories. His research interests include scientific computing, sequential and parallel algorithms for matrix computation, embedded computer systems, and compiler optimization of parallel programs. He received a PhD in computer science from Yale University. He is a member of SIAM and the ACM. Contact him at schreiber@hpl.hp.com.*

***B. Ramakrishna Rau** is an HP Fellow and director of the Compiler and Architecture Research Program at Hewlett-Packard Laboratories. His research interests include computer architecture, compilers, operating systems, and the automated design of computer systems. Rau received a PhD in electrical engineering from Stanford University. He is a Fellow of the IEEE and a member of the ACM. Contact him at rau@hpl.hp.com.*

***Darren C. Cronquist** is a research scientist in the Compiler and Architecture Research Program at Hewlett-Packard Laboratories. His research interests include reconfigurable computing, high-level synthesis, compilers, and embedded systems. He received a PhD in computer science from the University of Washington. He is a member of the IEEE and the ACM. Contact him at Darren_cronquist@hp.com.*

***Mukund Sivaraman** is a research scientist in the Compiler and Architecture Research Program at Hewlett-Packard Laboratories. His research interests include computer architecture, high-level synthesis, timing verification, delay-fault testing, and design automation of computer systems. He received a PhD in electrical and computer engineering from Carnegie Mellon University. He is a member of the IEEE and the ACM. Contact him at mukund@hpl.hp.com.*

Smart Cameras as Embedded Systems

Smart cameras capture high-level descriptions of a scene and perform real-time analysis of what they see. These low-cost, low-power systems push the design space in many dimensions, making them a leading-edge application for embedded system research.

Wayne Wolf
Burak Ozer
Tiehan Lv
Princeton University

Increasingly powerful integrated circuits are making an entire range of new applications possible. Complementary metal-oxide semiconductor (CMOS) sensors, for example, have made the digital camera a commonplace consumer item. These light-sensitive chips, positioned where film would normally be, capture images as reusable digital files that users can upload to their computer, manipulate with software, and distribute electronically.

Recent technological advances are enabling a new generation of *smart cameras* that represent a quantum leap in sophistication. While today's digital cameras capture images, smart cameras capture high-level descriptions of the scene and analyze what they see. These devices could support a wide variety of applications including human and animal detection, surveillance, motion analysis, and facial identification.

Video processing has an insatiable demand for real-time performance. Fortunately, Moore's law provides an increasing pool of available computing power to apply to real-time analysis. Smart cameras leverage very large-scale integration (VLSI) to provide such analysis in a low-cost, low-power system with substantial memory. Moving well beyond pixel processing and compression, these systems run a wide range of algorithms to extract meaning from streaming video.

Because they push the design space in so many dimensions, smart cameras are a leading-edge application for embedded system research. The Embedded Systems Group in Princeton University's Department of Electrical Engineering (http://www.ee.princeton.edu/~wolf/embedded-group/) has developed a first-generation smart camera system that can detect people and analyze their movement in real time.

DETECTION AND RECOGNITION ALGORITHMS

Although there are many approaches to real-time video analysis, we chose to focus initially on human gesture recognition—identifying whether a subject is walking, standing, waving his arms, and so on. Because much work remains to be done on this problem, we sought to design an embedded system that can incorporate future algorithms as well as use those we created exclusively for this application.

As Figure 1 shows, our algorithms use both low-level and high-level processing. The low-level component identifies different body parts and categorizes their movement in simple terms. The high-level component, which is application-dependent, uses this information to recognize each body part's action and the person's overall activity based on scenario parameters.

Low-level processing

The system captures images from the video input, which can be either uncompressed or compressed (MPEG and motion JPEG), and applies four different algorithms to detect and identify human body parts.

Region extraction. The first algorithm transforms the pixels of an image, like that shown in Figure 2a, into an $M \times N$ bitmap and eliminates the background. It then detects the body part's skin area using a YUV color model with chrominance values downsampled by a factor of two. Next, as Figure 2b illustrates, the algorithm hierarchically segments the frame into skin-tone and non-skin-tone regions by extracting foreground regions adjacent to detected skin areas and combining these segments in a meaningful way.

Contour following. The next step in the process, shown in Figure 2c, involves linking the separate

0018-9162/02/$17.00 © 2002 IEEE

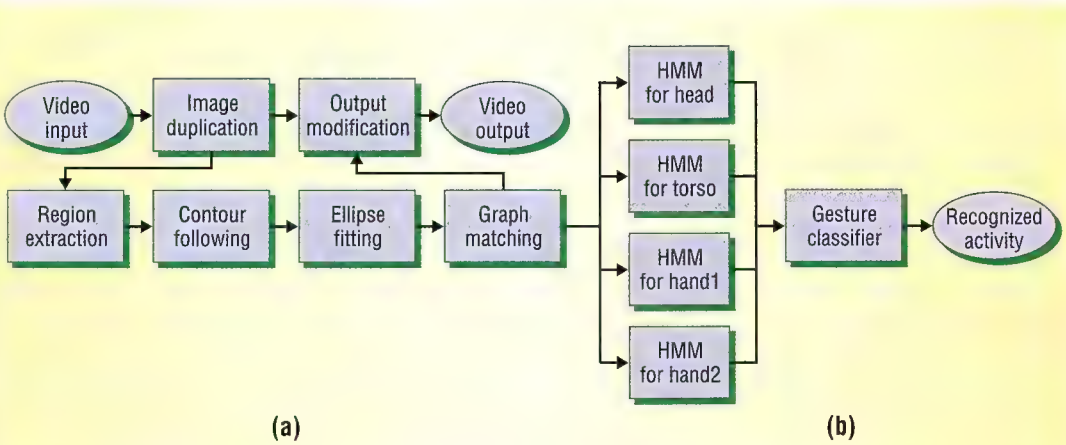

Figure 1. Human detection and activity recognition algorithms. (a) Low-level processing algorithms identify body parts and categorize their movements. (b) High-level processing algorithms use hidden Markov models (HMMs) and a gesture classifier to evaluate overall activity.

Figure 2. Initial steps in gesture recognition: (a) original image, (b) region extraction, (c) contour following, and (d) ellipse fitting.

groups of pixels into contours that geometrically define the regions. This algorithm uses a 3×3 filter to follow the edge of the component in any of eight different directions.

Ellipse fitting. To correct for deformations in image processing caused by clothing, objects in the frame, or some body parts blocking others, an algorithm fits ellipses to the pixel regions as Figure 2d shows to provide simplified part attributes. The algorithm uses these parametric surface approximations to compute geometric descriptors for segments such as area, compactness (circularity), weak perspective invariants, and spatial relationships.

Graph matching. Each extracted region modeled with ellipses corresponds to a node in a graphical representation of the human body. A piecewise quadratic Bayesian classifier uses the ellipses parameters to compute feature vectors consisting of binary and unary attributes. It then matches these attributes to feature vectors of body parts or meaningful combinations of parts that are computed offline. To expedite the branching process, the algorithm

Motion-Detection and Gesture-Recognition Systems

The research efforts focusing on human motion detection and gesture-recognition systems include Leonard,[1] a single-camera system that classifies simple motion events such as picking up an object using force dynamics. Mark Lucente, Gert-Jan Zwart, and Andrew D. George[2] have implemented a multimodal input system that also relies on one camera to let subjects manipulate virtual objects using gestures and voice commands. This system processes approximately 10 frames per second with a latency of 0.2 seconds.

The University of Maryland's Keck Laboratory for the Analysis of Visual Motion (http://www.umiacs.umd.edu/users/lsd/kecklab.html) employs a multicamera system to construct dynamic graphical representations of human movement and object manipulation. Digital cameras simultaneously capture some activity—such as a technician repairing a mechanism—from multiple viewpoints, and a suite of networked computers integrates this data with other sensor information into a 3D model for analysis using advanced computer graphics. Thomas B. Moeslund and Erik Granum discuss related work in their survey.[3]

Mircea Nicolescu and Gérard G. Medioni[4] have developed algorithms to electronically pan, tilt, and zoom through images supplied by an array of cameras. Their qualitative criteria for evaluating video input have confirmed that pan-tilt-zoom systems outperform wide-angle-lens cameras. Jonathan Foote and Don Kimber[5] have built a computationally and materially inexpensive panoramic camera system that also uses multiple cameras.

The MIT Media Lab (http://www.media.mit.edu/) is developing technology that can track people's actions, interpret gestures, and recognize facial expressions in environments ranging from the home and workplace to car interiors.[6-9] Smart rooms, smart desks, and wearable computers use context-sensing and communication devices to unobtrusively help people carry out everyday functions.

Scott Stillman and Irfan Essa[10] also have proposed a near-real-time system consisting of various types of sensors spread throughout an environment to track persons, detect faces, and recognize speech signals.

Other research efforts focus on various aspects of multiprocessor systems for video processing. Sek M. Chai and colleagues, for example, have developed an architecture for pixel-level processing in the imaging array.[11] In addition, John A. Watlington and V. Michael Bove are building a dynamically scheduled dataflow system using a distributed

begins with the face, which is generally easiest to detect.

High-level processing

The high-level processing component, which can be adapted to different applications, compares the motion pattern of each body part—described as a spatiotemporal sequence of feature vectors—in a set of frames to the patterns of known postures and gestures and then uses several hidden Markov models in parallel to evaluate the body's overall activity. We use discrete HMMs that can generate eight directional code words that check the up, down, left, right, and circular movement of each body part.

Human actions often involve a complex series of movements. We therefore combine each body part's motion pattern with the one immediately following it to generate a new pattern. Using dynamic programming, we calculate the probabilities for the original and combined patterns to identify what the person is doing. Gaps between gestures help indicate the beginning and end of discrete actions.

A quadratic Mahalanobis distance classifier combines HMM output with different weights to generate reference models for various gestures. For example, a pointing gesture could be recognized as a command to "go to the next slide" in a smart meeting room or "open the window" in a smart car, whereas a smart security camera might interpret the gesture as suspicious or threatening.

To help compensate for occlusion and other image-processing problems, we use two cameras set at a 90-degree angle to each other to capture the best view of the face and other key body parts. We can use high-level information acquired through one view to switch cameras to activate the recognition algorithms using the second camera. Certain actions, such as turning to face another direction or executing a predefined gesture, can also trigger the system to change views.

TOWARD AN EMBEDDED SYSTEM

As the "Motion-Detection and Gesture-Recognition Systems" sidebar describes, a number of researchers are working on human motion detection and gesture-recognition systems.

We initially used Matlab (http://www.mathworks.com/) to develop our algorithms. This technical computation and visualization programming environment runs orders of magnitude more slowly than embedded platform implementations, a speed difference that becomes critical when processing video in real time. We therefore ported our Matlab implementation to C code running on a very long instruction word (VLIW) video processor, which

resource manager for compact, relatively inexpensive media-processing applications.[12]

References

1. J.M. Siskind, "Visual Event Classification via Force Dynamics," *Proc. 17th Nat'l Conf. Artificial Intelligence* and *12th Conf. Innovative Applications of Artificial Intelligence* (AAAI/IAAI 00), AAAI Press/MIT Press, Menlo Park, Calif., 2000, pp. 149-155.
2. M. Lucente, G-J. Zwart, and A.D. George, "Visualization Space: A Testbed for Deviceless Multimodal User Interface," *Computer Graphics*, vol. 31, no. 2, 1997; http://www.lucente.biz/pubs/pubs.html.
3. T.B. Moeslund and E. Granum, "A Survey of Computer Vision-Based Human Motion Capture," *Computer Vision and Image Understanding*, vol. 81, no. 3, 2001, pp. 231-268.
4. M. Nicolescu and G.G. Medioni, "Electronic Pan-Tilt-Zoom: A Solution for Intelligent Room Systems," *Proc. IEEE Int'l Conf. Multimedia and Expo* (ICME 00), IEEE CS Press, Los Alamitos, Calif., 2000, pp. 1581-1584.
5. J. Foote and D. Kimber, "FlyCam: Practical Panoramic Video and Automatic Camera Control," *Proc. IEEE Int'l Conf. Multimedia and Expo* (ICME 00), IEEE CS Press, Los Alamitos, Calif., 2000, pp. 1419-1422.
6. A. Pentland and T. Choudhury, "Face Recognition for Smart Environments," *Computer*, Feb. 2000, pp. 50-55.
7. A.D. Wilson and A.F. Bobick, "Realtime Online Adaptive Gesture Recognition," *Proc. 15th Int'l Conf. Pattern Recognition* (ICPR 00), IEEE CS Press, Los Alamitos, Calif., 2000, pp. 270-275.
8. A. Pentland, "Smart Rooms, Smart Clothes," *Proc. 14th Int'l Conf. Pattern Recognition* (ICPR 98), IEEE CS Press, Los Alamitos, Calif., 1998, pp. 949-953.
9. A. Pentland, "Looking at People: Sensing for Ubiquitous and Wearable Computing," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 22, no. 1, 2000, pp. 107-119.
10. S. Stillman and I. Essa, "Towards Reliable Multimodal Sensing in Aware Environments," *Proc. Workshop on Perceptive User Interfaces* (PUI 01), ACM Press, New York, 2001; http://www.cs.ucsb.edu/PUI/PUIWorkshop/.
11. S.M. Chai et al., "Focal Plane Processing Architectures for Real-Time Hyperspectral Image Processing," *J. Applied Optics: Special Issue on Information Processing*, vol. 39, no. 5, 2000, pp. 835-849.
12. J. Watlington and V.M. Bove Jr., "A System for Parallel Media Processing," *Parallel Computing*, vol. 23, no. 12, 1997, pp. 1793-1809.

let us make many architectural measurements on the application and make the necessary optimizations to architect a custom VLSI smart camera.

Requirements

At the development stage, we evaluated the algorithms according to accuracy and other familiar standards. However, an embedded system has additional real-time requirements:

- *Frame rate.* The system must process a certain amount of frames per second to properly analyze motion and provide useful results. The algorithms we use as well as the platform's computational power determine the achievable frame rate, which can be extremely high in some systems.
- *Latency.* The amount of time it takes to produce a result for a frame is also important because smart cameras will likely be used in closed-loop control systems, where high latency makes it difficult to initiate events in a timely fashion based on action in the video field.

Moving to an embedded platform also meant that we had to conserve memory. Looking ahead to highly integrated smart cameras, we wanted to incorporate as little memory in the system as possible to save on both chip area and power consumption. Gratuitous use of memory also often points to inefficient implementation.

Components

Our development strategy called for leveraging off-the-shelf components to process video from a standard source in real time, debug algorithms and programs, and connect multiple smart cameras in a networked system. We use the 100-MHz Philips TriMedia TM-1300 as our video processor. This 32-bit fixed- and floating-point VLIW processor features a dedicated image coprocessor, a variable length decoder, an optimizing C/C++ compiler, integrated peripherals for concurrent real-time input/output, and a rich set of application library functions including MPEG, motion JPEG, and 2D text and graphics.

Our testbed architecture, shown in Figure 3, uses two TriMedia boards attached to a host PC for programming support. Each PCI bus board is connected to a Hi8 camera that provides NTSC composite video. Several boards can be plugged into a single computer for simultaneous video operations. The shared memory interface offers higher performance than the networks likely to be used in VLSI cameras, but they let us functionally implement and debug multiple-camera systems with real video data.

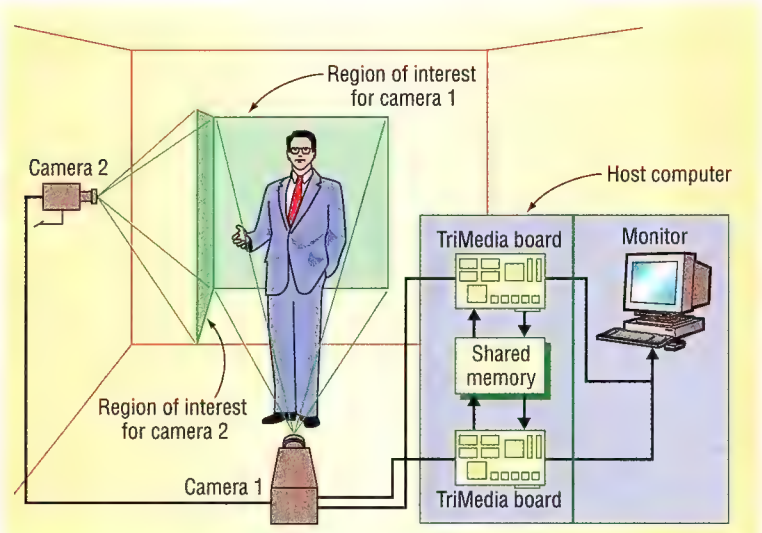

Figure 3. Smart camera test room and testbed architecture. Principal components include video processors on standard PCI bus cards, a shared memory interface, and a host PC for programming support.

EXPERIMENTS AND OPTIMIZATIONS

After converting the original Matlab implementation into C, we performed some experiments to gauge the smart camera system's effectiveness and evaluate bottlenecks. The unoptimized code took, on average, 20.4 million cycles to process one input frame, equal to a rate of 5 frames per second.

We first measured the CPU times of each low-level processing step to determine where the cycles were being spent. Microsoft Visual C++ is more suitable for this purpose than the TriMedia compiler because it can collect the running time of each function as well as its subfunctions' times.

Figure 4a shows the processing-time distribution of the four body-part-detection algorithms. Figure 4b shows the memory characteristics of each low-level processing stage.

As data representation becomes more abstract, input/output data volume decreases. The change in required memory size, however, is less predictable given the complex relationships that can form between abstract data. For example, using six single-precision, floating-point parameters to describe 100 ellipses requires only 2.4 Kbytes of memory, but it takes 10 Kbytes to store information about two adjoining ellipses.

Based on these early experiments, we optimized our smart camera implementation by applying techniques to speed up video operations such as substituting new algorithms better suited to real-time processing and using TriMedia library routines to replace C-level code.

Algorithmic changes

We originally fit superellipses (generalized ellipses) to contour points, and this was the most time-consuming step. Rather than trying to optimize the code, we decided to use a different algorithm. By replacing the original method developed from principal component analysis with moment-based initialization, we reduced the Levenberg-Marquardt fitting procedure, thus decreasing the execution time. Also, to accelerate processing during the graph-matching stage, we modified the algorithm to determine different regions' adjacency.

Library functions

During the region-extraction stage, the system processes each pixel in the input frame independently. Absolute value and threshold calculations result in branching, which limits instruction-level parallelism (ILP). One possible solution to this problem is to split the frame into several pieces and process these pieces on a multiprocessor or simultaneous multithreading platform. However, we opted to reduce the number of branches in the program.

The TriMedia processor provides an INONZERO operation that takes two input operands. If the first is not zero, the destination is set to the value of the second operand; otherwise, it is set to zero. Another special operation, IABS, can provide absolute values. These operations are visible in C code as they are packed into functions, and together they remove most of the branches.

We also used loop unrolling to extend basic block size. This optimization increased the processing speed of the region-extraction step by a factor of 2.3.

Control-to-data transformation

Increasing the processor's issue width can exploit the high degree of parallelism that region extraction offers. Using a processor with more functional units could thus reduce processing time during this stage. However, contour following, which converts pixels to abstract forms such as lines and ellipses, consumes even more time. The algorithm also operates serially: It finds a region's boundary by looking at a small window of pixels and sequentially moving around the contour; at each clockwise step it must evaluate where to locate the contour's next pixel. While this approach is correct and intuitive, it provides limited ILP.

We evaluated all possible directions in parallel and combined the true/false results into a byte, which served as an index to look up the boundary pixel in a table. We also manipulated the algorithm's control-flow structure to further increase ILP. These optimizations doubled the contour-following stage's running speed.

Optimization results

The combination of these methods radically improved CPU performance for the application. Optimization boosted the program's frame rate from 5 to 31 frames per second. In addition, latency

decreased from about 340 to 40-60 milliseconds per frame. We have since added HMMs and other high-level processing parts, and the program now runs at about 25 frames per second.

Our board-level system is a critical first step in the design of a highly integrated smart camera. Although the current system is directly useful for some applications, including security and medicine, a VLSI system will enable the development of high-volume, embedded computing products.

Because the digital processors and memory use advanced small-feature fabrication and the sensor requires relatively large pixels to efficiently collect light, it makes sense to design the system as two chips and house them in a multichip module. Separating the sensor and the processor also makes sense at the architectural level given the well-understood and simple interface between the sensor and the computation engine.

We believe that embedding single-instruction multiple-data (SIMD) processors into the sensor is not critical to achieve real-time performance. The advantages of leveraging existing sensor technology far outweigh any benefits of using pixel-plane processors until they become more plentiful. However, attaching special-purpose SIMD processors to the multiprocessor can be useful for boundary analysis and other operations. Such accelerators can also save power, which is important given the cost and effort required to deploy multiple cameras, especially in an outdoor setting. High-frame-rate cameras, which are useful for applications ranging from vibration analysis to machinery design, will likely require many specialized processing elements that are fast as well as area efficient, allowing the inclusion of more parallel units.

We are still in the early stages of determining the type of network best suited to a multicamera system. Distributed processing is an important means of reducing power consumption in such systems—sending raw pixels over the network is less efficient than sending the results of intermediate analysis. However, as algorithms for real-time multicamera analysis continue to develop, bandwidth requirements may change.

Acknowledgments

This work was funded by the New Jersey Center for Pervasive Information Technology; the MARCO/DARPA Center for Circuits, Systems, Software; and the National Science Foundation. We also thank Kees Vissers, Flaviu Turean, and Sebastian Mirolo at TriMedia, which generously provided us with processor boards, for their constructive suggestions.

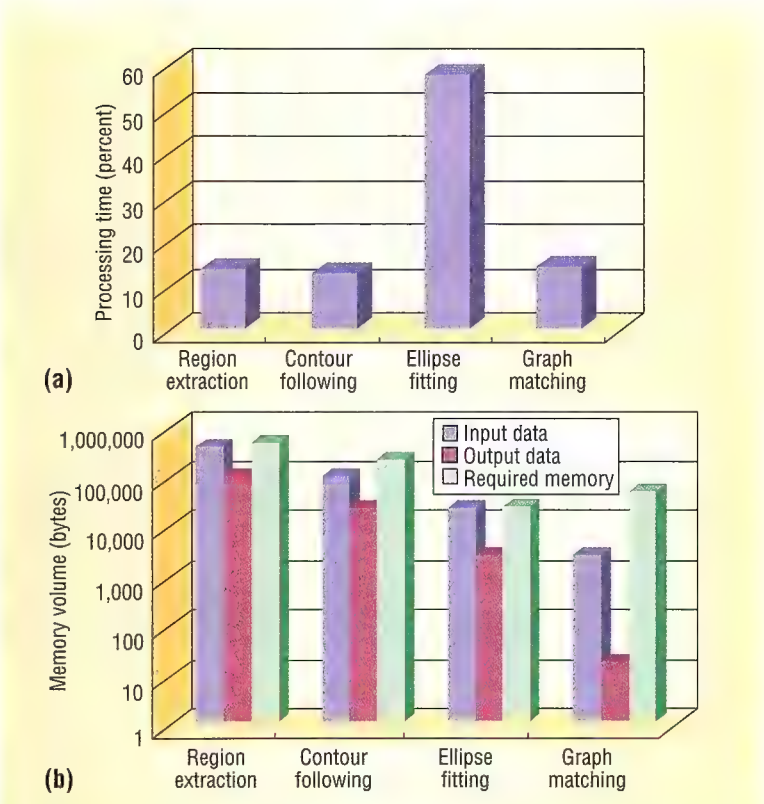

Figure 4. (a) Processing-time distribution of the original algorithms, implemented in C, before optimization. (b) Memory volume of low-level processing stages.

Wayne Wolf *is a professor in Princeton University's Department of Electrical Engineering, where he heads the Embedded Systems Group, and an associated faculty member in the Department of Computer Science. He received a PhD in electrical engineering from Stanford University. He is a Fellow of the IEEE and the ACM. Contact him at wolf@princeton.edu.*

Burak Ozer *is a research staff member and member of the Embedded Systems Group in the Department of Computer Engineering at Princeton University. He received a PhD in electrical engineering from the New Jersey Institute of Technology. He is a member of the IEEE. Contact him at iozer@ee.princeton.edu.*

Tiehan Lv *is a graduate student and member of the Embedded Systems Group in the Department of Electrical Engineering at Princeton University. He received an MAE in electrical engineering from Peking University, China. Contact him at lv@ee.princeton.edu.*

Mailbox-Based Scheme for Mobile Agent Communications

The authors present a flexible and adaptive scheme that associates each mobile agent with a mailbox but lets them decouple. Their 3D model provides a basis for evaluating existing communication protocols and allows for the design of new ones to meet various application requirements.

Jiannong Cao
The Hong Kong Polytechnic University

Xinyu Feng and Jian Lu
Nanjing University

Sajal K. Das
University of Texas at Arlington

Mobile agent technology has great potential for use in networking and distributed systems. Its applications range from telecommunications, e-commerce, and information searching to process coordination, mobile computing, and network management.[1] Mobile agents are autonomous objects or object clusters that move between locations in a mobile agent system—a distributed abstraction layer that provides the concepts and mechanisms for mobility and communication.[2,3]

Communication protocols are among the most important mechanisms in mobile agent systems.[4] In various situations, mobile agents at different hosts must cooperate with one another by sharing information and making decisions collectively.[5] To ensure effective interagent communication, these protocols must track target agent locations and deliver messages reliably.

In recent years, as the "Mobile Agent Tracking and Message Delivery Protocols" sidebar describes, researchers have proposed a wide range of schemes for agent tracking and reliable message delivery. However, each scheme has it own assumptions, design goals, and methodology. Protocol requirements with respect to different classes of applications are not well understood. As a result, no uniform or structured methods exist for characterizing current protocols, making it difficult to evaluate their relative effectiveness and performance.

In response to this problem, we propose a mailbox-based scheme for designing mobile agent communication protocols. This scheme assigns each agent a mailbox to buffer messages but decouples the agent and mailbox to allow them to reside at different hosts and migrate separately.

On the basis of the scheme's design space, we have developed a three-dimensional model that captures the main features of any communication protocol. This model provides a basis for evaluating various existing mobile agent communication protocols and for helping users design a flexible, adaptive protocol they can customize to meet specific application requirements.

PROTOCOL REQUIREMENTS

Communication protocols for mobile agents require location transparency, reliability, efficiency, asynchrony, and adaptability.

Location transparency

Because mobile agents can move autonomously from host to host, they cannot reliably "know" the locations of their communication peer. Therefore, a practical communication protocol must keep track of agent locations, allowing each agent to send messages to its peers without knowing where they physically reside.

0018-9162/02/$17.00 © 2002 IEEE

Researchers have recently proposed many mobile agent tracking protocols in different contexts ranging from mobile and wireless communications to wide-area distributed systems.[1] These protocols use various approaches that typically rely on some combination of a home server, forwarding pointers, broadcasts, and a hierarchical location directory.

Home server

Several mechanisms use a location server to keep track of a mobile object's current location. For example, the mobile Internet protocol uses a *home server* to route IP packets. A mobile host registers its care-of-address with its home host every time it moves. The home server directs all IP packets to the home host, which forwards them to the mobile one.

Mobile agent systems also use the home-server protocol as proposed in the Object Management Group's Mobile Agent System Interoperability Facility.[2] The protocol is simple to implement. Locating a mobile object incurs little communication overhead, but updating the locations and delivering messages incurs large costs.

To avoid a triangular routing problem, some researchers have proposed cache-based strategies.[3] The Internet mobile host protocol forwards packets along pointers left by the mobile host if a cache miss occurs.[4] However, these low-level protocols do not handle message loss due to host mobility.

Forwarding pointers

In some tracking mechanisms,[5] each host on a mobile agent's migration path keeps a *forwarding pointer* to the next host on the path. Each sender knows the target agent's home. Messages are sent to the agent's home and forwarded to the target object along the forwarding pointers. A path compression technique reduces message hops. After routing a message to the target object, the system sends an `Update_Entry` message back along the chain and updates forwarding pointers in the chain's nodes.

Other approaches

Many tracking protocols use broadcast and hierarchical approaches to implement distributed location management.

In a *broadcast* approach, the message sender broadcasts both location queries and pending messages to all system hosts.

However, simple broadcasts cannot avoid message loss that object mobility causes. A snapshot broadcast strategy[6] could be used to guarantee reliable message delivery to highly mobile agents as well as for group communications, but the large overhead makes broadcasts impractical in large-scale networks.

In the *hierarchical* approach,[7,8] a hierarchy of servers forms a location directory. The location database at each level contains location information for objects at levels below it. The hierarchy is usually tree-structured, with a leaf containing entries for all mobile objects in the corresponding unit zone. An internal node maintains data about mobile objects registered in its subtree's set of zones. For each object, the information is either a pointer to an entry at a lower-level location server or the object's actual current position.

Adaptive location management

A proposed *update* strategy for personal communication services optimizes location management cost on a per-user basis.[9] Each mobile user has an update strategy based on a unique mobility model and call arrival pattern, consisting of a set of binary decision variables—to update or not—for all areas. Users who enter a new location area can choose whether to update this information.

Different mobility patterns have prompted other update strategies based on timers, movement, distance, and state.[10]

Message delivery

Resending-based protocols[11] guarantee reliable message delivery to mobile agents. Using sliding-window mechanisms similar to those in the Internet transmission-control protocol, the sender can detect the loss of a message and resend it. After several resendings, the sender contacts the location manager and delivers the message to the receiver's new address. For highly mobile agents, there is no upper limit to the number of message resendings. Implementing reliable message delivery, however, requires synchronizing the message-passing operation with agent migration.

References

1. E. Pitoura and G. Samaras, "Locating Objects in Mobile Computing," *IEEE Trans. Knowledge and Data Eng.*, vol. 13, no. 4, 2001, pp. 571-592.
2. D. Milojicic et al., "MASIF: The OMG Mobile Agent System Interoperability Facility," *Proc. 2nd Int'l Workshop Mobile Agents* (MA 98), Springer-Verlag, Berlin, 1998, pp. 50-67.
3. Y. Lin, "Determining the User Locations for Personal Communications Services Networks," *IEEE Trans. Vehicular Technology*, vol. 43, no. 3, 1994, pp. 466-473.
4. C. Perkins, A. Myles, and D.B. Johnson, "IMHP: A Mobile Host Protocol for the Internet," *Computer Networks and ISDN Systems*, vol. 27, no. 3, 1994, pp. 479-491.
5. J. Desbiens, M. Lavoie, and F. Renaud, "Communication and Tracking Infrastructure of a Mobile Agent System," *Proc. 31st Ann. Hawaii Int'l Conf. System Sciences* (HICSS 98), IEEE CS Press, Los Alamitos, Calif., 1998, pp. 54-63.
6. A.L. Murphy and G.P. Picco, "Reliable Communication for Highly Mobile Agents," *Proc. 1st Int'l Symp. Agent Systems and Applications and 3rd Int'l Symp. Mobile Agents* (ASA/MA 99), IEEE CS Press, Los Alamitos, Calif., 1999, pp. 141-150.
7. M. van Steen et al., "Locating Objects in Wide-Area Systems," *IEEE Comm. Magazine*, vol. 36, no. 1, 1998, pp. 104-109.
8. W.V. Belle, K. Verelst, and T. D'Hondt, "Location Transparent Routing in Mobile Agent Systems—Merging Name Lookups with Routing," *Proc. 7th IEEE Workshop Future Trends of Distributed Computing Systems* (FTDCS 99), IEEE CS Press, Los Alamitos, Calif., 1999, pp. 207-212.
9. S.K. Sen, A. Bhattacharya, and S.K. Das, "A Selective Location Update Strategy for PCS Users," *ACM/Baltzer J. Wireless Networks*, vol. 5, no. 5, 1999, pp. 311-326.
10. V.W-S. Wong and V.C.M. Leung, "Location Management for Next-Generation Personal Communication Networks," *IEEE Network*, vol. 14, no. 5, 2000, pp. 18-24.
11. M. Ranganathan, M. Bednarek, and D. Montgomery, "A Reliable Message Delivery Protocol for Mobile Agents," *Proc. Joint 2nd Int'l Symp. Agent Systems and Applications and 4th Int'l Symp. Mobile Agents* (ASA/MA 00), Springer-Verlag, Heidelberg, Germany, 2000, pp. 206-220.

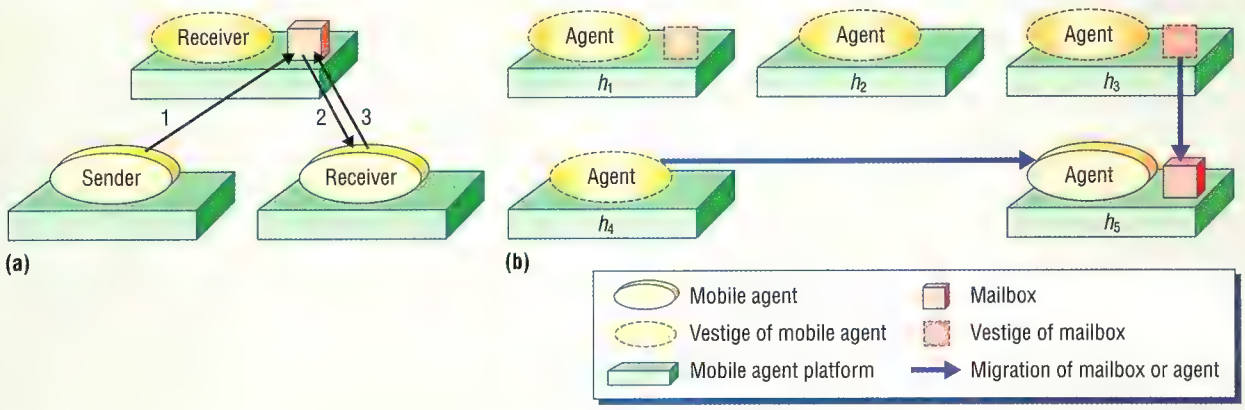

Figure 1. Mailbox-based scheme. (a) An agent sends a message to another agent's mailbox (1), from which the receiver later obtains it via either a (2) push or (3) pull operation. (b) A mobile agent's mailbox migrates at a lower frequency than the agent itself.

Reliability

The asynchronous message passing and agent migration in mobile agent systems can result in message loss. Even an ideal fault-free transport mechanism cannot guarantee the delivery of messages to their destination agents. A practical protocol, however, must ensure that all messages are routed to the target agent, no matter how frequently it migrates, in a bounded number of hops.

Efficiency

Protocol costs include the number of messages sent, their size, and the distance traveled. An efficient protocol minimizes these costs while supporting two operations: agent migrations to new sites and message delivery, which includes first locating the target agent. However, minimizing the costs of these operations leads to conflicting requirements.

Consider the tradeoff between a *full-information* strategy in which every site in the network maintains complete up-to-date information about every agent's location, and a *no-information* strategy that does not update agent migration information.[6] Delivery is inexpensive with the former approach, but updating location information at every site makes migration expensive. With the latter strategy, on the other hand, migration costs nothing, but delivery requires searching the entire network.

Any protocol must balance these costs either generally or within some specific communication and migration pattern.

Asynchrony

Although a communication protocol should coordinate message forwarding with agent migration to guarantee reliable message delivery, it should not overly constrain agent mobility by frequent and tight synchronization. Agents should be able to freely migrate to other hosts whenever necessary.

In addition, the agent's autonomous and asynchronous execution should not rely heavily on the agent's home to locate and deliver every message to them. Agents should be independent of the process that created them, and the home site should be able to disconnect from an agent as soon as it migrates.

To retain the advantages of asynchrony in mobile agent systems, any communication protocol must address asynchronous requirements in both agent migration and execution.

Adaptability

Different applications often have their own requirements. Some, for example, favor asynchrony, while reliability is more important in others. Various interagent communication and agent migration patterns also may have unique implications on migration and delivery strategies. Protocols can be designed for specific applications to achieve optimal performance, but an adaptive protocol that can suit as many kinds of applications as possible in a general-purpose mobile agent system is more desirable.

MAILBOX-BASED SCHEME

In our scheme, each mobile agent has a mailbox that buffers the messages sent to it. As Figure 1a shows, an agent can direct a message to another agent's mailbox, and the receiving agent uses a push or pull operation to obtain the message from the mailbox.

Although logically part of the agent, the mailbox can be detached from its owner—the agent can leave the mailbox at one host as it migrates to a new one. Interagent communication thus consists of two distinct steps:

- transmission of a message from the sender to the receiver's mailbox, and
- delivery of the message from the mailbox to its owner agent.

Mailboxes are mobile objects, but because they are not autonomous, they cannot determine their migration path. Any existing message delivery strategy will satisfy the first step, but because mailboxes migrate at a much lower frequency than their associated agents, any protocol design must include a

parameter for when to move the mailboxes.

Figure 1b shows an agent migrating sequentially from host h_1 to h_5 but taking its mailbox only while moving to h_1, h_3, and h_5. By definition, the set of hosts on the mailbox's migration path is a subset of the hosts on the mailbox owner's migration path. The mobile agent's home is the first host on the migration paths of both the agent and its mailbox.

Because it is rare for two mobile agents roaming the Internet to use synchronous communication, we assume that mobile agent communication is largely asynchronous. The Internet's lengthy and unpredictable message delays, which can easily last several seconds, likewise prohibit frequent use of synchronous communication in a mobile agent application.

Our scheme is built on top of a reliable network communication layer, which guarantees that messages will not be lost during transmission and will be delivered between hosts.

DESIGN OPTIONS

To meet an application's specific requirements, our scheme offers choices in three aspects of protocol design: mailbox migration frequency, mailbox-to-agent message delivery, and synchronization of message forwarding with object migration.

Mailbox migration frequency

The number of mailbox migrations during a mobile agent's life cycle, along with the times when the migrations occur, can vary.

No migration (NM). A mobile agent can move alone, leaving its mailbox at home during the agent's life cycle. All messages are sent to the home, and the agent uses a mailbox-to-agent delivery mode to obtain messages. Tracking the mailbox carries no cost, but the agent's home must forward all messages. This triangular routing[7] increases the communication overhead for message delivery.

Full migration (FM). As part of the mobile agent's data, the mailbox continuously migrates with the agent. Although the cost of message delivery between mailbox and agent is zero, tracking the mailbox is difficult. Frequent agent-mailbox migration necessitates a tradeoff between message loss and the cost of guaranteeing message delivery.

Jump migration (JM). A mobile agent can determine dynamically whether to take its mailbox before each migration. In making the decision, the agent may consider the number of messages it will receive at its target host, the distance between the target host and the host where its mailbox currently resides, and other factors.

An agent that seldom receives messages at its target host doesn't need to take its mailbox to the new host. However, if an agent expects to receive messages frequently and its target host is far from the host where its mailbox currently resides, fetching messages from the remote mailbox will be expensive. In this case, the agent should migrate with its mailbox to the target host.

In JM mode, a protocol can work more flexibly when based on a decision suited to particular agent migration and interagent communication patterns, reducing the cost of both tracking and delivery operations.

The Internet's unpredictable message delays prohibit frequent use of synchronous communication in a mobile agent application.

Mailbox-to-agent message delivery

Messages destined to an agent are all sent to the agent's mailbox, and the agent later receives the messages by either a push or pull operation.

Push (PS). The mailbox keeps its owner agent's address and forwards every message to it. Although message queries incur no costs, the agent must notify the mailbox of its current location after every migration. If the agent migrates frequently but communicates with other agents at only a few hosts on the migration path, most message delivery location registration messages would be superfluous and introduce a significant migration overhead. PS mode is needed if real-time message delivery is required.

Pull (PL). The mobile agent retrieves messages from its mailbox's address whenever needed. The mailbox doesn't need to know the agent's current location, thus avoiding location registration, but the agent must query its mailbox for messages. Polling messages would increase message delivery overhead.

Migration-delivery synchronization

Our scheme lets users determine whether they need reliable message delivery. If users require high reliability, they can overcome message loss by

- synchronizing the *host*'s message forwarding and the *mailbox*'s migration (SHM),
- synchronizing the *mailbox*'s message forwarding and the *agent*'s migration (SMA), or
- both, known as full synchronization (FS).

NS denotes the extreme case of no synchronization performed.

Figure 2 shows how synchronization occurs between the message-forwarding object (mobile agent server or mailbox) and the moving object

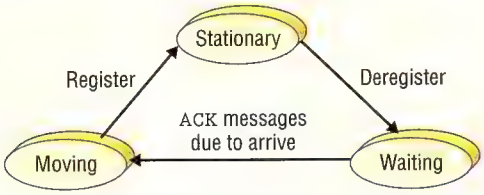

Figure 2. State switching of a mobile object. During synchronization, the communication protocol can forward messages to the mobile object in either a stationary or waiting state, but it must block messages when the object is moving.

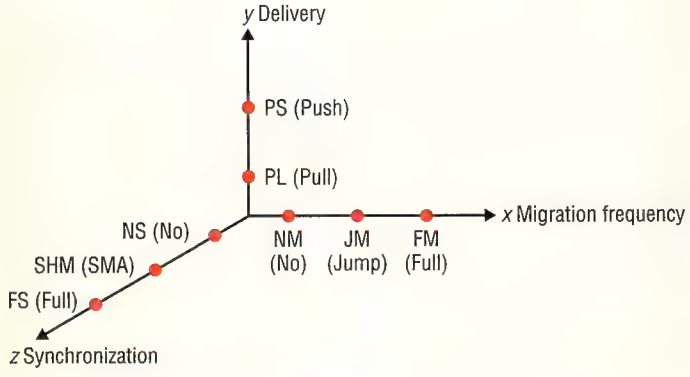

Figure 3. Three-dimensional design space. Each axis represents one design aspect and contains a range of properties that designers can combine with properties from other dimensions in various ways.

Table 1. Parameter combinations and corresponding protocols.

Protocol	Location registration	Reliability
NM-PS-NS	Yes (agent → mailbox)	No
NM-PS-SMA	Yes (agent → mailbox)	Yes
NM-PL-NS	No	Yes
FM-*-NS	No	No
JM-PL-NS	No	No
JM-PS-NS	Yes (agent → mailbox)	No
FM-*-SHM	Yes (mailbox → host)	Yes
JM-PL-SHM	Yes (mailbox → host)	Yes
JM-PS-FS	Yes (agent → mailbox, mailbox → host)	Yes

(mailbox or mobile agent). Before migrating, the moving object sends `Deregister` messages to all objects that might forward messages to it and waits for `ACK` messages from each forwarding object, which contains the number of messages to be forwarded. Once the messages arrive, the moving object continues on its path. After migration, the moving object informs all message-forwarding objects of its arrival by sending them `Register` messages.

3D DESIGN MODEL

These three design options generate a three-dimensional design space. As Figure 3 shows, each option represents one orthogonal dimension with a constraint spectrum. Because the three dimensions are independent of one another, designers can combine properties from different dimensions in various ways. The full range of properties for a given application can thus vary greatly.

Combining parameters from all three dimensions yields a taxonomy of mobile agent communication protocols. A string of the format XX-YY-ZZ expresses a protocol in which XX represents mailbox migration frequency (NM, JM, or FM), YY stands for mailbox-to-agent message delivery (PL or PS), and ZZ symbolizes synchronization (NS, SHM, SMA, or FS). A protocol's overall configuration has a special value for each of the three parameters. Most combinations have plausible applications.

Table 1 shows the different protocols derived from our mailbox-based scheme, with the description of their location registration modes and whether they can satisfy the required reliability. An asterisk denotes that multiple values are applicable.

Home-server protocols

All NM-mode protocols adopt a home-server approach. In this case, the agent's home acts as the message-forwarding server.

NM-PS-NS is identical to mobile IP.[7] The agent registers its current location with its mailbox residing at its home. Messages are sent to the mailbox, which pushes them to its owner agent. This protocol does not guarantee message delivery: If the agent migrates during message forwarding, the message will be lost.

Reliable message delivery requires synchronizing agent migration and message forwarding from the mailbox. This results in NM-PS-SMA, a synchronized version of mobile IP. In NM-PL-NS, the agent pulls messages from its mailbox, ensuring message delivery without using synchronization.

Home-server protocols are simple and work well for small-to-medium systems with relatively few agents. However, triangular routing increases communication overhead, especially when agents are widely distributed. In a system with numerous agents and frequent interagent communication, the home may become a performance bottleneck and a single point of failure. Further, mobile agent dependence on the home as a message-forwarding server constrains asynchronous execution.

Forwarding-pointer-based protocols

In FM-*-NS, each host on the mailbox migration path keeps a forwarding pointer to the successive host. The sender caches the target mailbox's previously obtained location; if the cache has no such

address, the sender uses the target agent's home as the mailbox's cached address. Messages are sent to the cached address directly; if a cache miss occurs, the messages are forwarded along the pointers.

When the mailbox receives the message and discovers that the sender has outdated knowledge of its address, the mailbox notifies the sender of its current location, and the sender updates the cached address. Because the mailbox is bound with its owner agent in the FM mode, no remote interaction occurs between the mailbox and agent.

JM-PL-NS and JM-PS-NS are similar to FM-*-NS except that the mailbox migrates in JM mode, which is a kind of path compression technique. When the agent and its mailbox are at different hosts, the agent uses the pull and push operations to get messages from its mailbox.

The forwarding-pointer scheme has no location update cost. Because the sender delivers messages to the target agent's cached address, the agent's home workload is smaller. Even if the cache is outdated, messages can be routed to the target agent along the path. If one host on the migration path fails, however, the target agent is no longer reachable. More importantly, this protocol cannot guarantee message delivery because a message can keep chasing the target agent if the agent migrates frequently.

The FM-*-NS protocol's multihop path could degrade communication performance significantly. In contrast, an agent's mailbox migrates less frequently in JM-PL-NS and JM-PS-NS, thus shortening the message-forwarding path and reducing the communication overhead. The chasing problem is also less likely to occur in these two protocols.

Distributed-registration-based protocols

The last three protocols shown in Table 1 use synchronization to guarantee reliable message delivery. Before migrating, the mailbox informs all hosts on its migration path and waits for ACK messages from them. After arriving at the target host, the mailbox registers its new address with hosts on its migration path.

The sender sends messages to the cached mailbox address—say, h_k. If the mailbox has moved, h_k forwards the messages to the mailbox's current address and notifies the sender of the new address. This protocol is similar to NM-PS-SMA, but it distributes the agent home's role to all hosts on the migration path. It can also be regarded as a forwarding-pointer protocol with migration-based path compression—the agent updates all pointers on its migration path after one or several migrations.

If an agent migrates frequently, synchronization and location registration would make FM-*-SHM prohibitively expensive. However, if the mailbox migrates in JM mode, both the mailbox registration times and the number of hosts on its migration path decrease.

In JM-PL-SHM, the agent uses the pull operation to obtain messages from its mailbox. To guarantee message delivery, the protocol must synchronize host message forwarding and mailbox migration. In JM-PS-FS, the mailbox pushes each incoming message to its owner agent, so the protocol must also synchronize message pushing from the mailbox and the owner agent's migration. For this reason, the protocol uses full synchronization.

To guarantee message delivery, the protocol must synchronize host message forwarding and mailbox migration.

In a recent analysis of JM-PL-SHM,[8] we found synchronization to be quite effective. Messages are forwarded at most once to reach the receiver's mailbox, and no chasing problem exists. By properly determining the mailbox's migration frequency, designers could create a protocol that reduces both migration and delivery costs.

Our proposed scheme offers only a guideline for designing and classifying mobile agent communication protocols. The 3D model is open to further detail along each dimension to support uniform definitions of the protocol terms, assumptions, and properties. Our future work will address the design of adaptive communication protocols, moving beyond static combinations of the three design options and giving application programmers the dynamic capabilities to customize protocols at runtime. ■

Acknowledgments

We are grateful to the anonymous referees for their helpful suggestions. Our research is partially supported by the University Grant Council of Hong Kong under CERG grant B-Q518 (PolyU 5076/01E), the Hong Kong Polytechnic University under PolyU ICRG grant A-PC53, and the US National Science Foundation under grants EIA-0086260 and EIA-0115885.

References

1. D.B. Lange and M. Oshima, "Seven Good Reasons for Mobile Agents," *Comm. ACM*, vol. 42, no. 3, 1999, pp. 88-89.
2. M. Strasser, J. Baumann, and F. Hohl, "Mole—A

Java Based Mobile Agent System," *Special Issues in Object-Oriented Programming: Workshop Reader of the 10th European Conf. Object-Oriented Programming* (ECOOP 96), M. Mühlhaüser, ed., Dpunkt Verlag, Heidelberg, Germany, 1996, pp. 327-334.
3. V.A. Pham and A. Karmouch, "Mobile Software Agents: An Overview," *IEEE Comm. Magazine*, vol. 36, no. 7, 1998, pp. 26-37.
4. P.T. Wojciechowski, *Algorithms for Location-Independent Communication between Mobile Agents*, tech. report DSC-2001/13, Communication Systems Dept., Swiss Federal Inst. of Technology, Zurich, Switzerland, 2001.
5. J. Cao et al., "Checkpointing and Rollback of Wide-Area Distributed Applications Using Mobile Agents," *Proc. 15th Int'l Parallel and Distributed Processing Symp.* (IPDPS 01), CD-ROM, IEEE CS Press, Los Alamitos, Calif., 2001.
6. B. Awerbuch and D. Peleg, "Online Tracking of Mobile Users," *J. ACM*, vol. 42, no. 5, 1995, pp. 1021-1058.
7. C. Perkins, ed., *IP Mobility Support*, Internet Engineering Task Force RFC2002 (Standards Track), 1996; http://www.ietf.org/rfc/rfc2002.txt.
8. X. Feng et al., "An Efficient Mailbox-Based Algorithm for Message Delivery in Mobile Agent Systems," *Proc. 5th IEEE Conf. Mobile Agents* (MA 01), IEEE CS Press, Los Alamitos, Calif., 2001, pp. 135-151.

Jiannong Cao *is an associate professor in the Department of Computing at the Hong Kong Polytechnic University, Hung Hom, Kowloon, Hong Kong. His research interests include parallel and distributed computing, computer networks, Internet computing, mobile computing, and fault tolerance. Cao received a PhD in computer science from Washington State University. He is a member of the IEEE Computer Society, the ACM, and the American Association for the Advancement of Science. Contact him at csjcao@comp.polyu.edu.hk.*

Xinyu Feng *is a master's student in the Department of Computer Science and Technology at Nanjing University, Nanjing, China. His research focuses on networking, distributed systems, and mobile agents. Feng received a BSc in computer science from Nanjing University.*

Jian Lu *is a professor in the Department of Computer Science and Technology at Nanjing University. His research interests include Internet computing, mobile agents, componentware, and parallel object-oriented technology. Lu received a PhD in computer science from Nanjing University. He is a member of the ACM. Contact him at lj@nju.edu.cn.*

Sajal K. Das *is a professor in the Department of Computer Science and Engineering and founding director of the Center for Research in Wireless Mobility and Networking (CReWMaN) at the University of Texas at Arlington. His research interests include resource and mobility management in wireless networks, mobile and pervasive computing, mobile Internet architectures and protocols, distributed processing, and grid computing. Das received a PhD in computer science from the University of Central Florida, Orlando. He is a member of the IEEE Computer Society, the IEEE, the ACM, the New York Academy of Sciences, and Sigma Xi. Contact him at das@cse.uta.edu.*

Tough times make it more important than ever to cut risks and increase the return on your technology investments. COMDEX Fall 2002 is your best opportunity to experience, explore, and learn about the latest product innovations, leading vendors, and powerful new business strategies. With expanded Educational Programs, special events, and visionary Keynote speakers, including Bill Gates, Carly Fiorina, and Brian Halla, COMDEX gives you the insight you need to take your organization further.

- FREE access to the Exhibit Floor
- FREE admission to the Keynotes and SuperSession
- FREE entrance to the Great Debates and Hot Spots
- FREE admission to the new Digital Lifestyles Conference
- Pricing discounts (through September 1) on Educational Programs

Don't miss the Business Technology Conference featuring **Crystal Ball: Technology Infusions.**

Presented in partnership with IEEE SPECTRUM

Register today! Visit www.comdex.com to register or call 888-508-7510

COMDEX EDUCATIONAL PROGRAMS: NOVEMBER 16–21, 2002
COMDEX EXHIBITION: NOVEMBER 18–22, 2002
LAS VEGAS CONVENTION CENTER, LAS VEGAS, NEVADA
OFFICIAL CORPORATE SPONSORS OF KEY3MEDIA GROUP

Official Card of COMDEX

Mercedes-Benz

NYSE New York Stock Exchange

Key3Media Group PRESENTS COMDEX FALL 2002

THE GLOBAL TECHNOLOGY MARKETPLACE

Click and Vote Online!

http://computer.org/election/

IEEE COMPUTER SOCIETY

IEEE Computer Society Election

Nominees for IEEE Computer Society Office and Board of Governors

On the following pages are the position statements and biographies of the IEEE Computer Society's candidates for president-elect, first and second vice presidents, and Board of Governors. Within each category, candidates are listed in alphabetical order. Election of officers to one-year terms and of Board members to three-year terms, each beginning 1 January 2003, will be by vote of the membership as specified in the bylaws.

Ballots must be returned no later than 12:00 noon EDT on Friday, 4 October 2002. Members in all regions can vote via the Web at http://computer.org/election/ or by fax to election.com inc. at +1 516 248 7790. Members without Web access in region 8 should return ballots to the Brussels office; in region 10, to the Tokyo office. All others should return ballots to the IEEE Computer Society, c/o election.com inc., PO Box 9217, Garden City, NY 11530-9217, USA. For replacement ballots, call +1 866 720 4357. Results will be announced in the December issue of *Computer*.

The opinions expressed in the statements are those of the individual candidates and do not necessarily reflect IEEE Computer Society positions or policies.

Nominees for president-elect

Carl K. Chang

Position Statement. During my sustained service to the IEEE Computer Society, my VISA—Vision, Interoperability, Strategy, and Action—concept has prevailed.

Vision. As the leading provider of technical information and services to the world's computing professionals, the Society should strengthen its capability to offer members sharply focused, on-demand learning opportunities to stay current and do so cost-effectively through the Internet. The Society should also continue its global leadership role in both emerging technology and burning professional and policy issues.

Interoperability. Once the vision is recognized and shared, we must build a sound infrastructure to foster interoperability. We need effective operational interoperability between volunteers and Society staff, communicative interoperability among members around the world, and technical interoperability among various segments comprising the computing profession. In addition, the interoperability between the central IEEE and the Society must be revamped to halt the erosion of financial stability that has come about in recent years.

Strategy. The Society must assess and redevelop its strategic plan based on the collective wisdom of its volunteer leaders and Society staff, with your input. The effectiveness, however, of executing such plans must be rooted in a common vision and executed with great interoperability.

Action. We as a Society must choose and act strategically on the most important tasks now. Among others, the following tasks must take highest priority:

- Work with the IEEE to grant the Society leadership more financial and operational autonomy to better serve its own members.
- Work with computing professionals and Society members to voice our concerns and offer our expertise to help with challenges such as counter-terrorism.
- Guide the Society to assemble, develop, and disseminate technical information and lifelong-learning materials in various forms that are timely, focused, relevant, and of high quality, so that we may deliver genuine value to our members worldwide.

My leadership and experience as a volunteer since 1979 (international conferences, technical activities, press and publications, and education) provide me with a solid background to champion these tasks. My employer recognizes the significance of such an honorable

(Chang continued on next page)

Gerald L. Engel

Position Statement. It is an honor to be a candidate for IEEE Computer Society president. At this time, our profession—and our Society—face significant opportunities and significant challenges. I believe I can offer effective leadership to the Computer Society as together we face those opportunities and challenges.

Looking ahead, it is vital that the Society consolidate and strengthen existing programs of publications, conferences, and services, buttressed by an effective partnership of volunteers and staff. At the same time, we must accomplish three important goals: strengthen and expand our developing programs that provide electronic access to our dynamically developing field, provide expanded continuing education opportunities to our membership that will enhance our skills, and further develop and enhance our quality assurance programs for both academic programs and individual practitioners.

We must effectively broaden our membership to be still more inclusive, such that Society membership reflects and represents all areas of information technology. We must also be more proactive in expanding our activities worldwide, assuring all IT practitioners that they have a home in the IEEE Computer Society.

It is equally important that our relationship with other organizations, both inside and outside the IEEE, be strengthened and extended. To be effective, we must continue to identify new partners who will join us in better serving our communities throughout the world.

Within the IEEE, we must continue to enhance our leadership position and work to make appropriate changes that ensure the future success of both the Society and the IEEE. I have effectively worked in the past with all the candidates for 2003 IEEE president-elect and believe this should lead to opportunities for improving our relationships with the IEEE.

Finally, we must continue to seek ways to make your membership more valuable and useful. I will look to you as members to help me identify new areas and approaches we have not used in the past and to strengthen those that we know do work. I will always be accessible to work with you on your ideas.

I look forward to working with you and the Computer Society. Thank you for considering my candidacy.

(Engel continued on next page)

(Chang continued)

service to the profession and has granted me sufficient release time. I will be honored to serve you. Please feel free to contact me at c.chang@computer.org to advise me or discuss any concern that you might have.

Biography. Vice president for Educational Activities since 2001, Carl Chang served previously as vice president for Press Activities (1999) and secretary of the Board of Governors (1998). He has been a member of the Planning, Audit, and Awards Committees and has contributed actively on the Educational Activities, the Publications, and the Conferences and Tutorials boards. He was editor in chief of *IEEE Software* (1991-1994) and, in that capacity, established an Industry Advisory Board. He organized the 1997 President's Roundtable, commissioned to answer practitioners' needs. He spearheaded the joint Society and ACM task force on Computing Curricula 2001 for education renovation. He has cofounded and chaired several IEEE and non-IEEE international conferences. His sustained contributions have been recognized with the Society's Outstanding Contribution Award, Meritorious Service Award, Golden Core recognition, and the IEEE Third Millennium Medal.

Chang received a PhD in computer science from Northwestern University. After starting his career in industry at GTE Automatic Electric and Bell Laboratories, Chang joined the University of Illinois at Chicago in 1984. In July, he became chair for the Department of Computer Science at Iowa State University in Ames. He has published extensively in both software engineering and net-centric computing. Chang is an IEEE Fellow.

(Engel continued)

Biography. Gerald Engel is the 2001-2002 vice president for conferences and tutorials. Active in the Society for 25 years, he has served as vice president for education, among other positions. He is a past president of the IEEE Society on the Social Implications of Technology. In addition, Engel is president of the Computer Science Accreditation Board (CSAB) and a Board of Directors member of the Accreditation Board for Engineering and Technology (ABET). He has also served as an IEEE Division Director.

Engel is the Leonhardt Professor of Computer Science and Engineering at the University of Connecticut Stamford Campus. He holds a doctorate from Pennsylvania State University. Prior to coming to Connecticut, he held teaching and research positions at Old Dominion University, Christopher Newport College, and the Virginia Institute of Marine Science. From 1991 to 1995, he was a program director in the National Science Foundation's Computer and Information Science and Engineering Directorate, acting as deputy division director of Computer and Computation Research for 1994-1995. His research interests have focused on computer science and engineering education and on social and ethical issues in the field. He has received numerous awards from the Society, including the 2000 Computer Society Merwin Distinguished Volunteer Service Award.

Nominees for first vice president

Lowell G. Johnson

Position Statement. PLEASE VOTE, even if it is not for me (but I would appreciate your vote). I'd rather have 50 percent of the ballots returned than actually win this election, but that's wishful thinking, since historically the return rate has been 15-18 percent.

However, I would still like to hear members' concerns or suggestions on how to improve the Society. I prefer e-mail at L.Johnson@computer.org, or phone +1 651 635 7305 during business hours. My thanks to those who responded last year.

The senior volunteers and staff continue to work on all the standard issues, including finances and politics with the IEEE, staffing, Web upgrades, publications (paper and electronic), intellectual property rights, and many others. But we are also working on new initiatives to both provide more services and allow the members to do more for themselves. One of our biggest problems is the dissemination of information about new tools or services. How do you reach people who do not regularly visit our Web sites or do not want to receive mass e-mailings?

One of our best recent successes is the distance learning program. Members can now take roughly 100 free courses electronically or visit some courses many times without completing them, thus using the HTML course, for example, as reference material. This is a terrific deal for members, and we are working to add more capability and sources for future years. Please give this a try.

Thank you for your support, and I encourage you to contact me with comments or suggestions.

Biography. Lowell Johnson, the 2002 second vice president for Standards Activities, is a member of the IEEE Standards Board, the Standards Association Corporate Advisory Group, and NesCom (the New Standards Committee), which he chairs. He chairs the Portable Applications Standards Committee (Posix), the Standards Activities Board, and is the Society's liaison to the IEEE Standards Board. Johnson has chaired or been a member of numerous Society committees and has chaired several standards working groups. He is a senior IEEE member.

A consultant at Unisys where he has worked for more than 25 years, Johnson received a BS in mathematics and an MS in physics

(Johnson continued on next page)

Deborah K. Scherrer

Position Statement. Our Society's vitality and strength come from its ability to recognize and meet the needs of its worldwide and professionally diverse members. We all look to the Society for:

- access to top-quality technical information,
- opportunities for continuing education,
- professional support such as standards and curriculum development, and
- forums in which to share ideas and interrelate with colleagues.

The Society must be nimble and proactive in new initiatives. We need to respond to the increasing diversity of computing professionals and adapt our publications, conferences, activities, and standards accordingly. We need to expand our services in continuing education, improve the quality and timeliness of our publications and conferences, enhance training opportunities for computing professionals, and work harder at becoming a catalyst for university-industry coordination. Most importantly, we need to find new models to trigger and support special-interest communities and provide them the services they require.

Changes in IEEE attitude and policies continue to have significant impacts on our organization. So far, we have succeeded in protecting member programs from the changes. During this transition, the Society requires a strong, proactive, and courageous leadership able to contend with restructuring, able to communicate impacts with members, and capable of assuring the unhampered startup and continuation of dynamic Society programs. I care deeply about the Society and will use whatever strengths I can bring to assist in this transition.

It would be an honor to continue as your first vice president. I pledge my time and energy toward supporting our Society.

Biography. In addition to serving as the Society's first vice president, Deborah Scherrer chairs the Technical Activities Board. Her service includes the Executive Committee, Board of Governors, Education Activities Board, and Planning Committee. Previously, she was Society secretary; coordinator of the Technical Activities Forum in *Computer*; member of the Conferences and Tutorials Board and the Awards, Audit, Membership, Finance, and Nominations

(Scherrer continued on next page)

(Johnson continued)

from the University of Minnesota. While completing his PhD thesis in high-energy physics, he nearly lost his life in an auto accident with a drunk driver, which eventually forced him to leave the university and begin his career in computer science rather than in physics. Johnson has received Certificates of Appreciation from the IEEE, the Society, and the US Department of Commerce. He has received several Working Group Chair Awards, the Distinguished Service Award, and the IEEE Third Millennium Medal. His Computer Society honors include the Meritorious Service Award, the Distinguished Service Award, and Golden Core recognition.

(Scherrer continued)

Committees; and Society ombudsman. She represents the Society as an alternate on the Computing Sciences Accreditation Board.

Scherrer was Usenix Association president and on its Board of Directors for 11 years. She has actively served on numerous conference committees and participated on the editorial boards of *Computer, Computing Systems, Unix Review,* and *Unix World*.

Currently working on NASA programs at Stanford University, Scherrer previously managed the Automation Group at Transmeta. She has been president of Mt. Xinu, staff scientist at Lawrence Berkeley Laboratory, and scientific programmer at the Crimean Astrophysical Observatory. She also founded her own PC software company. Scherrer completed her undergraduate and graduate work at UC Berkeley.

Scherrer received the Usenix 1996 Lifetime Achievement Award, a Unix Systems Laboratory "Academic Driver" award, the Society Golden Core recognition, and the IEEE Third Millennium Medal. She is an honorary lifetime member of Usenix.

Nominees for second vice president

Wolfgang K. Giloi

Position Statement. In my career, I have worked as a computer engineer and company executive in industry and as a professor of computer science at universities in Germany and the United States. For 15 years, I have directed a computer research institute that was involved in many international research programs in Europe, Japan, and the US. This background has given me a profound understanding of the needs of computer scientists and engineers in industry and at universities worldwide. As a member of the IEEE Computer Society's Board of Governors and as chair of the Audit Committee, I initiated the Computer Society's Central and Eastern Europe Activities Program, as well as introducing a change of its bylaws that strengthens the representation of non-US areas in the Society's governing bodies.

The Computer Society is acting in a rapidly changing world. Our challenge is to nimbly adapt to these changes, to remain the provider of quality services to our members, who have made us the largest society within the IEEE. In recent years, the biggest challenge has been the transition from paper to electronic publishing.

We must not rest on our laurels. We have launched new initiatives that will make us a total information provider for our members as well as serve as an institution for continuing education through distance learning.

Having served in past years in administrative functions, it is my desire to utilize my professional expertise by focusing, as second vice president, on these technical activities.

Biography. Wolfgang Giloi has been on the Board of Governors of the Computer Society for 11 years. He chaired the Audit Committee and served on various other committees. In 2001, he was Society secretary; currently, he is Society treasurer. He was a development engineer and the general manager of the Analog/Hybrid Computer Department of AEG-Telefunken, professor of electrical engineering at the Technical University of Berlin (1965-1971), computer science professor at the University of Minnesota (1972-1977), and research professor and director of the Institute for Computer Architecture and Software Technology of the German National Research Center for Information Technology (1978-1997). He is a member of the Berlin-Brandenburg Academy of Science and the German Convent of Technical Sciences. He served as adjunct professor of computer science at UCLA (since 1986) and is an honorary professor at Shanghai Jiao Tong University.

Giloi was involved in national and international research programs (Suprenum, Genesis, and RWC). He designed the Suprenum supercomputer, which from 1989 to 1991 was the most powerful supercomputer worldwide. An IEEE Fellow, he has also received the Ring of Honor of the Verein Deutscher Ingenieure, the IFIP Silver Core, the Computer Society Golden Core recognition, and the Rudolf-Diesel-Medal.

Rangachar Kasturi

Position Statement. Publication of periodicals with high-quality, peer-reviewed content is perhaps the most significant service the IEEE Computer Society offers its members. During the past several years, I have worked closely with volunteer leaders and with our publications staff to maintain excellence in quality of content and to significantly improve publication timeliness and relevance.

Through my efforts and the support of Society leadership, the IEEE recently withdrew the changes it had made to its copyright policies to include references to the Digital Millennium Copyright Act and the US Export Control Regulations. These changes had severely restricted our ability to serve our members and the public at large through dissemination of open, public scientific knowledge.

The Society provides member services through six vice-president-led programs (chapters, conferences and tutorials, education, publications, standards, and technical activities). Our services are far reaching and wide ranging—although this is our strength, special efforts are sometimes required to facilitate board interaction and cooperation because of the large organizational structure. With my experience as an active participant on several of these program boards, I can provide effective leadership for increased cooperation among the boards.

All IEEE societies are facing significant challenges during these times of financial strain at the parent organization. We must provide leadership and be innovative in our approach to ensure that the vitality of our organization remains strong and the service to our members is not curtailed. As second vice president and a voting member of the Society's Board of Governors, I would support such activities.

Biography. Rangachar Kasturi has served as the vice president for publications since 2001. He served as the general co-chair for the IEEE Conference on Computer Vision and Pattern Recognition (2001) and as a member of the Fellow Evaluation Committee (1999-2001). He has served as the Transactions Operations Committee chair (2000), as editor in chief of *IEEE Transactions on Pattern Analysis and Machine Intelligence* (1995-1998), and as a lecturer in the Distinguished Visitor Program (1987-1990).

A professor of computer science and engineering and of electrical engineering at Penn State University since 1982, Kasturi previously worked as a communication systems design engineer for about 10 years. He is the coauthor of the textbook, *Machine Vision* (McGraw Hill), and of three tutorial texts on computer vision and document image analysis published by the Computer Society.

Kasturi received an MSEE (1980) and a PhD (1982) from Texas Tech University and a BE in electrical engineering (1968) from Bangalore University. A Fellow of the IEEE, he has been honored with a Fulbright Fellowship, an IEEE Computer Society Meritorious Service Award, the Society's Golden Core recognition, a Penn State Engineering Premier Researcher Award, and a membership in the Texas Tech Electrical Engineering Academy.

Board of Governors nominees (11 nominees; vote for seven)

Oscar N. Garcia

Position Statement. After many years of IEEE Computer Society involvement, it is good to be back to face the issues confronting us today. I see three major relationships we could improve:

With our members. Improved quality of publications and meetings, distance education and certification, chapter activities, the digital library, and other benefits could enhance our standing with members.

With the profession. Leadership on technical, educational, and socioeconomic issues, and improved relations with sister societies, would benefit the profession.

With the IEEE. As the largest member Society, we could lead on standards, professional, and regional activities, and also help to solve some of the financial problems that face the IEEE as a whole and therefore, by association, face us.

Much has changed since my days as Society president, and the Society has grown in complexity and stature. I will be honored to continue my service to this outstanding organization if you were to elect me.

Biography. Oscar Garcia served as Computer Society president for almost two and a half years and also in many other capacities. He has also served the IEEE as a Board of Directors member. An IEEE and AAAS Fellow, he has received the Society's Merwin Award and the IEEE's Emberson Award, both for highest service to the profession. He is currently the NCR Distinguished Professor at Wright State University and was director of the Interactive Systems Program at the National Science Foundation's Computer and Information Science and Engineering Directorate. Prior to this, Garcia was full professor of electrical engineering and computer science at George Washington University and was charter chair of the Department of Computer Science and Engineering at the University of South Florida.

Garcia's recent research has dealt with topics in complexity, bioinformatics, human-computer interaction, artificial intelligence, expert systems, and software engineering. He developed courses for IBM, popular in the US and Japan, on artificial intelligence and expert systems. His earlier areas of research were robust speech recognition, computer architecture and parallel processing, testing of digital circuits, and arithmetic coding theory. Garcia received BS and MS degrees from North Carolina State University and a PhD from the University of Maryland.

Mark A. Grant

Position Statement. If re-elected, I will continue to work to keep the IEEE Computer Society strong, financially secure, and independent. As a member of the Audit Committee, I will continue to demand justification for expenses allocated to us by our parent organization. My business background offers a unique perspective in audit matters, and I will continue to fight those who would take our hard-earned monies.

The Board of Governors has benefited from my legal perspective for challenges to issues that went unchallenged in the past. This will continue to be a significant advantage for the Board when confronted with legal justifications for actions imposed on the Society that are unjustifiable under any logical standard.

Finally, it is an advantage for any team to have diversity of skills. If re-elected, I will continue to focus my unique blend of legal, business, and technical skills on the issues most critical to the future of the Computer Society.

Biography. Mark Grant is a member of the 2001-2002 Board of Governors, a member of the Audit Committee, and a Society Golden Core member.

A practicing attorney specializing in patent law, Grant has 30 years' experience in semiconductor technology, electronics, computers, software, and related fields. He has held legal positions such as vice president and general counsel at Mosel Vitelic and as director of intellectual property at National Semiconductor. He has also held engineering and management positions such as microprocessor specialist at Texas Instruments, project manager at Chrysler, and design engineer at General Motors. Grant's current legal practice includes litigation strategy and management, intellectual property counseling, technical product analysis, technology licensing, and preparation and coordination of design programs to avoid infringement.

Grant received a BS in electrical engineering with honors from Oakland University, an MS in biomedical engineering from the University of Michigan, and a Juris Doctor, cum laude, from Santa Clara University School of Law. He is admitted to practice in California, the District of Columbia, before the US Court of Appeals for the Federal Circuit, before the US Patent and Trademark Office, and before the US Tax Court.

Michel Israel

Position Statement. As an international organization, the IEEE Computer Society should be of benefit to its members worldwide. This can be done by developing cooperation with national technical societies. In these times of global exchanges, we must develop our educational activities in continuing education and international student exchange. This is particularly important for computer science, which is practiced in an international environment.

Most major computer systems and software are developed for worldwide use. It is, therefore, all the more important for computer scientists to understand the language and culture of major trading partners. If elected, I will use my international experience to

- promote information exchange between the Society and national professional societies by means of Web discussion forums,
- provide more educational opportunities by making the delivery of distance learning courses more efficient and more readily available, and
- facilitate the development of curricula for mutual recognition of accredited programs throughout the world.

Biography. Michel Israel's 20 years of Society service include vice president for Technical Activities (2000), chair of the Central and Eastern European Initiative Committee (CEEIC) for 1999-2000, treasurer (1998-1999), Design Automation Technical Committee chair (1994-1998), secretary (1997), ombudsman, and chair of the European Activities Committee. He served on the Nomination Committee (2001) and on the Publications, Membership, and Conferences and Tutorials boards. Israel has been the coordinator for the new French Association on Information Technology and has served on France's Ministry of Education Accreditation Board for Computer Science.

Israel is the scientific counsellor of the French Embassy in Tokyo and is developing a Japan-France doctoral college. At the University of Evry, he was dean of the Faculty of Sciences and chair of the CNRS computer science laboratory. He was the EU chair of an EU-US exchange program (1997-2000) and a visiting professor at Turkey's University of Galatasaray (1999).

Israel, who earned a PhD in computer sciences from Paris 6 University, has numerous Society awards, including Distinguished Service as vice president of the Technical Activities Board and as CEEIC chair, and Outstanding Contribution for establishing the USSR's first computer chapter. He is a Golden Core member and received the IEEE Third Millennium Medal.

David G. McKendry

Position Statement. We must continue to improve the overall value of IEEE Computer Society membership for all members. Every member should be able to list at least three important benefits they value enough to remain a member. We must reassess the benefits we offer as they relate to the benefits that members value most.

We must improve industry support of volunteerism by practitioners by demonstrating the added value of membership for their employees. Participation in activities at the chapter level as well as on boards and committees must become more acceptable and even encouraged by industry leaders and managers.

We must increase the involvement of members, especially younger professionals. This involvement starts with chapter officer experience to develop volunteers with an understanding of what is required to start or run a chapter. These experienced chapter officers should then be encouraged to move into our international boards and committees.

Biography. David McKendry serves on the Chapter Activities Board and the Technical Activities Board. Previously, he served on the Board of Governors (1999-2001), as Division VIII PACE coordinator (1994-2000), as assistant vice president of the Membership Development Committee (1992-1998), and on the Standards Activities Board (1998), Membership Activities Board (1991-1998), and Area Activities Board (1988-1990). He also chaired the Distinguished Visitors Program (1993-1994) and led the rewrite of DVP policies and procedures. His longtime service includes chairing the Southeast Michigan Chapter (1986-1993)—reviving an inactive chapter that later won the Society's first Outstanding Chapter Award. McKendry also chaired the Southeast Michigan IEEE Section in 1996.

McKendry is a staff analysis engineer for General Motors, where he has worked since 1980. Previously, he worked for General Electric Medical Systems and Carnation, and he also served in the US Air Force. He is a US patent holder and author of several publications.

He earned an MBA from Wayne State University (1988), and a BSEE (1980) and MSEE (1983) from Marquette University. He was awarded Golden Core recognition in 1996 and the IEEE Third Millennium Medal in 2000. He is a member of the Society of Automotive Engineers and a Senior Member of the Computer Society.

Sorel Reisman

Position Statement. In my 17+ years with the IEEE Computer Society, I've participated in many of our magazines' developmental stages, from inception to production—as author (*CG&A*); contributing writer (*Computer*); columnist (*IEEE Software*); editorial board member (*Software, IEEE MultiMedia, IEEE IT Professional*); associate editor in chief (*ITPro*); task force and committee member creating new magazines, portals, and electronic communities; EIC Search Committee chair; Publications Board member; and as Magazine Operations Committee chair. I've seen major changes in our demographics and participated in efforts to develop new products and services for our traditional constituency while providing value for new, diverse members. My participation in the Electronic Products and Services Committee is an important way to continue developing new programs necessary to sustain our growth. My experience will be invaluable on the Board of Governors, in helping to develop electronic communities, redesign the Society's Web site, and work more closely with the IEEE to realize mutual objectives.

Biography. Sorel Reisman's Society activities include serving on the editorial boards of *IEEE Software, IEEE MultiMedia,* and *IT Professional*. He has been a member of the Publications Board and served two terms as chair of the Magazine Operations Committee. He is a member of the Electronic Products and Services Committee.

Reisman has been a professor of information systems in the College of Business at the California State University, Fullerton, since 1986, where he is also an Academic Technology coordinator. He serves as Director of Systemwide Academic Technology Services at the California State University Office of the Chancellor. His industrial experience includes management positions at IBM, Toshiba America, and Thorn EMI. He received a BASc (electrical), an MA, and a PhD from the University of Toronto.

Reisman teaches information systems courses in e-commerce, information systems management, and information systems development. His research interests have focused on multimedia computing, information systems management, and, more recently, electronic communities and distance learning systems. He has presented and published numerous articles on information systems in professional journals. He is the author of the book, *Multimedia Computing: Preparing for the 21st Century*, and is working on a new book on best practices in electronic learning communities.

Stephen B. Seidman

Position Statement. If I am elected to the Board, my primary concern will be to maintain the IEEE Computer Society's financial integrity and independence. At the same time, my goal for the Society's programs will be to ensure that we continue to support high-quality education for students and for working professionals. This goal is based on the Society's critical mission to support education and workforce development for the world's computing professionals.

The Society has long supported curriculum development and undergraduate program accreditation in computer science, information systems, computer engineering, and software engineering. In addition, the Society offers software professionals a wide range of educational opportunities, from conference tutorials to online courses. This year has seen a Society-sponsored certification program introduced for software development professionals, which entails an examination as well as face-to-face and online exam preparation courses. As a member of the Board, I would foster, encourage, and extend all of these activities.

Biography. Stephen Seidman is a member of several Computer Society volunteer organizations: the Professional Practices Committee, the Educational Activities Board, and the Technical Activities Board. He has played a leading role in the Society's effort to develop a certification examination for software engineers, and he currently chairs the committee responsible for developing training materials for the examination.

In 2001, Seidman became the founding dean of the College of Computing Sciences at the New Jersey Institute of Technology. He was chair of the Department of Computer Science at Colorado State University from 1996 to 2001 and head of the Department of Computer Science and Engineering at Auburn University from 1990 to 1996. From 1972 to 1990, Seidman was a faculty member at George Mason University, where he was a founding member of the Department of Computer Science.

Seidman received a BS in mathematics from the City College of New York, and an MA and a PhD in mathematics from the University of Michigan. His research interests are in software architectures, and he has published more than 50 technical papers. Seidman is a recipient of the IEEE Third Millennium Medal.

Kathleen M. Swigger

Position Statement. From 1995 to 1998, I served as one of the software editors for *Computer* magazine. I was then appointed to *Computer*'s executive board, and currently serve the magazine as Special Issues editor. I also serve on the Conferences and Tutorials Board, Audit Committee, Distance Learning Committee, and Electronic Products and Services Committee, where I am working on a special project to redesign the IEEE Computer Society's Web site.

If elected, I would continue to explore ways in which technology can be used to provide communities for our members. I am also concerned about how technology is creating an impermanent work environment and, as a consequence, what the Society can do to ensure that its members will cope with and even profit from these changes. Finally, I will explore ways in which technology can be used to give the international membership more access to and representation on the Society's boards and committees.

Biography. Kathleen Swigger is a computer science professor at the University of North Texas in Denton. Her research and teaching involve computer-supported cooperative work, human interfaces, and artificial intelligence, specifically intelligent interfaces. Currently, Swigger is working with the faculty of the Middle East Technical University in Ankara, Turkey, on a project to examine the cultural factors that contribute to the success or failure of groups of programmers collaborating via the Internet. She is also involved with several technology projects related to undergraduates and women.

Swigger has substantial practical experience including contributing to the development of an intelligent tutoring system for space operations, a gate allocation system for a major airline, a Pilot Mission Debriefer for the Air Force, and several cooperative software projects with different companies.

As a professor, Swigger has earned a number of teaching and professional awards. In addition, she has received several NSF, US Department of Education, US Department of Defense, and state grants to develop special collaborative software. She earned a BA, an MA, and a PhD from the University of Iowa.

Makoto Takizawa

Position Statement. As the leading professional society in the computing and networking fields, the IEEE Computer Society's role in the 21st century will be crucial. The Society should promote both academic and industry professionals with international conferences, publications, education, forums, and standardization. The global collaboration of practitioners in different countries, communities, and organizations is essential.

If elected, I want to improve—and enhance—the Society's ability to promote international technical activities, especially involving members outside North America. In the future, half of the Society's members will come from outside North America; roughly 40 percent do now. I would like to help the local Computer Society chapters in these countries become more involved and foster their cooperation with Computer Society chapters in the US. I also plan to promote technical activities such as conferences, forums, and publications by means of high-speed Internet technologies.

Biography. Makoto Takizawa founded the IEEE International Conference on Information Networking and the IEEE International Conference on Parallel and Distributed Systems in Asia. He served as program co-chair of the IEEE International Conference on Distributed Computing Systems in 1998 and as program vice chair of ICDCS in 1994 and 2000. Currently, he serves as a general co-chair of ICDCS 2002 and of IEEE ISORC.

Takizawa is a professor in the Department of Computers and Systems Engineering and a dean of the graduate school of Science and Engineering at Tokyo Denki University in Japan. From 1998 to 2002, he chaired the Information Division at the university's Research Institute for Technology. Previously, he was a visiting professor at GMD-IPSI, Germany (1989-1990), and he has been a visiting professor at Keele University, England, since 1990.

A fellow of the Information Processing Society of Japan, Takizawa served on the IPSJ executive board from 1998 to 2000. He was also an IPSJ *Journals* editor from 1994 to 1998.

Takizawa received BE, ME, and DE degrees in computer science from Tohoku University, Japan. In 1996, he won the best-paper award at ICPADS. He is a member of the IEEE, the ACM, and IPSJ.

Michael R. Williams

Position Statement. I am privileged to have served the IEEE Computer Society in several areas, including a one-year replacement position on the Board of Governors. This experience showed me that the Society's major challenge is gaining control of the financial situation caused by the central IEEE having seen fit to spend our reserves to cover its deficits. This problem, if it continues, threatens us, and our sister societies, with bankruptcy. Confronting this problem must be the Computer Society's first priority, and I will focus my efforts on its solution.

While it is necessary to control the damage being done to our reserve funds, we must not forget the fact that our members are interested in the benefits they receive as part of their membership—particularly the publications. I will continue to use my experience in the publications area to promote these benefits and ensure that they are kept up to the highest standards.

Biography. Michael Williams has been involved in Society activities for many years, serving this past year as a member of the Board of Governors. He served first as a member of the editorial board for *IEEE Annals of the History of Computing*, spent several years as assistant editor in chief, and eventually served two terms as editor in chief. He was a member of the Publications Board for the past seven years, chaired the Society's History Committee, chaired the Pioneer Awards Committee, and is now completing a term as editor in chief of the Computer Society Press, during which time he has also served as liaison to the IEEE Press.

For 30 years, Williams was a professor of computer science at the University of Calgary, but he recently left that post to become head curator at the Computer History Museum in Mountain View, California, where he is endeavoring to create a world-class museum to preserve our industry's heritage.

Williams holds a BSc in chemistry and a PhD in computer science. He has received several awards for Computer Society service, a national award from the Canadian Information Processing Society honoring his work in recording computer history, and University of Calgary recognition for teaching contributions.

Zhiwei Xu

Position Statement. As the world's largest society in our profession, the IEEE Computer Society should continue to take a leadership role in the 21st century, serving as a catalyst for change. The Society's vision is to be the leading provider of technical information and services to the world's computing professionals. If elected, I would strive to realize this vision by

- enhancing the Society's connectivity and service-providing channels through technology and social engineering means;
- increasing the Society's visibility and relevance among the world's computing professionals, especially outside North America;
- improving the Society's education services to the world;
- responding to the needs of our volunteers and members, to take advantage of the constantly changing nature of our profession; and
- promoting the increasing involvement of international membership in the Society's activities, encompassing education, conferences, publications, chapters, boards, and committees.

Biography. Zhiwei Xu currently serves as representative of the Computer Society Beijing Center and as chair of the IEEE Beijing Section (2001-2002). He also served on the Society's Publications Board (1999-2001) and Task Force on Cluster Computing (2000-2002). He helped establish the Beijing Center and wrote two books to promote the computer profession and the Computer Society. He is editor in chief of the *Journal of Computer Research and Development* and chair of the International Cooperation Committee of China Computer Federation.

Xu is a professor and deputy director of the Institute of Computing Technology, Chinese Academy of Sciences, where he conducts research in high-performance computing and grid technology. He led the development of China's first national computational grid and was chief architect of Dawning superservers, which, as reported in *Science*, were instrumental in helping bioscientists to discover the draft sequence of the rice genome.

Xu holds an MS from Purdue University (1984) and a PhD from the University of Southern California (1987). He is a senior member of the IEEE and a member of the ACM and the Global Grid Forum. He received several national prizes for his research work and a Best Teacher award from the Chinese Academy of Sciences.

Yervant Zorian

Position Statement. With today's trend of rapidly advancing information technologies, including progress in wireless, networking, and communications, the challenge to each computing professional is not only to keep pace with, but also to contribute to, this advancement.

Our Society has traditionally played a key role in the professional development of its members to prepare them to meet such challenges. However, these state-of-the-art technologies have made accomplishing this goal very difficult. If elected to the Board, I plan to contribute to our Society's vitality by encouraging members to adapt to today's changes and ensuring that our offerings remain the prime source of information to our members, whether they are in academia or industry. At the same time, I would support efforts to provide these offerings through electronic media such as online tutorials, interactive workshops, electronic publications, and the like. Such electronic forums leverage our traditional top-quality offerings and provide much wider penetration and global access.

Biography. Yervant Zorian has actively served the IEEE Computer Society for the past 16 years in various capacities. He is currently the Society's vice chair for communities on the Technical Activities Board and is also the editor in chief emeritus of *IEEE Design & Test of Computers*. He is a past chair of the Test Technology Technical Council. Under his leadership, that Council saw major growth in membership and offerings, and hence was elevated from a Technical Committee to a Technical Council.

Zorian founded and has served as the chair of several Society activities, including the IEEE P1500 standardization Working Group, the IEEE Workshop on Testing Embedded Core-Based System-on-Chips, and, recently, the IEEE Workshop on Test Resource Partitioning. He participates on the editorial advisory boards of *IEEE Spectrum* and *JETTA*.

Zorian is vice president and chief scientist of Virage Logic and chief technology advisor of LogicVision. He received an MSc from the University of Southern California and a PhD from McGill University.

Zorian has published three books, received several best paper awards, and holds 10 US patents. He is a Society Golden Core member, an Honorary Doctor of the National Academy of Sciences in Armenia, and a recipient of the Society's Distinguished Services Award. He is also a Fellow of the IEEE.

IEEE Computer Society Election

Cast your vote quickly and easily on the Web at http://computer.org/election/ or via fax to +1 516 248 7790. Ballots must be received no later than 12:00 noon EDT on Friday, 4 October 2002. To vote by mail, use the return-mail envelope provided and send your ballot to the address for your region.

Regions 1-7 and 9:
IEEE Computer Society
c/o election.com inc., PO Box 9217
Garden City, NY 11530-9217, USA

Region 8:
European Office, 13, Avenue de l'Aquilon
B-1200 Brussels, Belgium

Region 10:
Asia/Pacific Office, Watanabe Building
1-4-2 Minami-Aoyama
Minato-ku, Tokyo 107-0062, Japan

PURPOSE The IEEE Computer Society is the world's largest association of computing professionals, and is the leading provider of technical information in the field.

MEMBERSHIP Members receive the monthly magazine **COMPUTER**, discounts, and opportunities to serve (all activities are led by volunteer members). Membership is open to all IEEE members, affiliate society members, and others interested in the computer field.

COMPUTER SOCIETY WEB SITE The IEEE Computer Society's Web site, at **http://computer.org**, offers information and samples from the society's publications and conferences, as well as a broad range of information about technical committees, standards, student activities, and more.

OMBUDSMAN Members experiencing problems—magazine delivery, membership status, or unresolved complaints—may write to the ombudsman at the Publications Office or send and e-mail to help@computer.org.

CHAPTERS Regular and student chapters worldwide provide the opportunity to interact with colleagues, hear technical experts, and serve the local professional community.

IEEE COMPUTER SOCIETY

AVAILABLE INFORMATION

To obtain more information on any of the following, contact the Publications Office:

- Membership applications
- Publications catalog
- Draft standards and order forms
- Technical committee list
- Technical committee application
- Chapter start-up procedures
- Student scholarship information
- Volunteer leaders/staff directory
- IEEE senior member grade application (requires 10 years practice and significant performance in five of those 10)

To check membership status or report a change of address, call the IEEE toll-free number, +1 800 678 4333. Direct all other Computer Society-related questions to the Publications Office.

PUBLICATIONS AND ACTIVITIES

Computer. An authoritative, easy-to-read magazine containing tutorial and in-depth articles on topics across the computer field, plus news, conferences, calendar, industry trends, and product reviews.

Periodicals. The society publishes 12 magazines and 10 research transactions. Refer to membership application or request information as noted at left.

Conference Proceedings, Tutorial Texts, Standards Documents. The Computer Society Press publishes more than 160 titles every year.

Standards Working Groups. More than 200 groups produce IEEE standards used throughout the industrial world.

Technical Committees. Thirty TCs publish newsletters, provide interaction with peers in specialty areas, and directly influence standards, conferences, and education.

Conferences/Education. The society holds about 100 conferences each year and sponsors many educational activities, including computing science accreditation.

IEEE President-Elect Candidates Address Computer Society Concerns

Within the computing field, the IEEE Computer Society has established a reputation for excellence. As a component of the IEEE, the Computer Society's activities parallel those of 40 other societies and councils serving the engineering and computing disciplines. Representing by far the largest IEEE society contingent, the Computer Society has more than 92,400 members, approximately 74 percent of whom are full IEEE members.

Recognizing the influence and control the IEEE wields over our Society and in turn the power of Society members' votes to influence the IEEE leadership, we posed five questions to this year's candidates for IEEE president-elect. Because this election determines who will serve as president-elect in 2003, president in 2004, and past president in 2005—vital positions within the IEEE's governing body—our members must cast informed votes.

Our volunteer leaders have identified the following questions as essential to the Society, the IEEE, and the Society's relationship with the IEEE. The first response to each question states the Computer Society position. These positions synthesize the views of our most senior leadership: the Society's current, past, and incoming presidents, and the division directors who represent the Society on IEEE governing bodies.

We present these questions and answers to help you make your decision in the IEEE election, which closes 1 November. We also remind and encourage you to cast your vote for Computer Society leaders by 4 October in our Society election.

—Stephen L. Diamond, IEEE Computer Society president-elect

Question

1 Vibrant societies are essential to the IEEE's spirit of innovation and the introduction of new products and member services. In the past, societies funded new products and services with their own financial reserves. Recently, societies lost effective control over use of their liquid reserves, and the IEEE adopted a system of central control over and approval of all new initiatives valued over $50,000.

As president, how do you see the IEEE balancing the need for minimal, prudent financial controls with the need for organizational nimbleness and society empowerment?

Computer Society Position

1 Many IEEE society leaders believe that, because the reserves they built have been redirected to cover continuing central IEEE deficits, they are losing the ability to develop new products. The approval process that new IEEE business rules impose may delay the launching of products and services and cause the societies to miss opportunities or lose their competitive edge in providing top services to professionals.

One way to balance financial controls with organizational empowerment is to delegate decision-making authority to responsible societies, allowing those with proven track records to operate within agreed-to guidelines. Societies that have a history of prudent financial management should be made accountable for their strategic decisions and projects. Up to a certain threshold, they would report new initiatives to the IEEE, rather than request approval.

The limit of $50,000 for new initiatives is too low for the larger societies; the size of the society's budget should be factored into this limit.

Vijay K. Bhargava is president of Binary Communications and holder of the Canada Research chair in Broadband Wireless Communication at the University of Victoria. His IEEE experience includes leadership in technical, regional, and educational activities and as president of the IEEE Information Theory Society. He received a PhD in 1974 from Queen's University.

For information on his candidacy for IEEE president-elect, visit http://www.ece.uvic.ca/~bhargava/ieee.

Luis T. Gandía is the founder and president of L. Gandía & Associates, Inc., a 40-year-old manufacturer's representative firm serving the electrical and telecommunication industries in the Caribbean. His IEEE experience includes leadership roles in the IEEE Board of Directors as vice president, secretary, Division VI director, and regional director.

For information on his candidacy for IEEE president-elect, visit http://www.luistgandia.com.

Arthur W. Winston is an IEEE Life Fellow with extensive IEEE experience up to and including vice president. He is a member of several IEEE societies, including the Computer Society. His 40 years of industrial, academic, and senior management experience include responsibility for the Temperature Measurement System for the Apollo spacecraft. Winston is currently the director of the Gordon Institute of Tufts University.

For information on his candidacy for IEEE president-elect, visit http://www.arthurwinston.com.

Vijay K. Bhargava

The IEEE Computer Society is a large society with excellent management infrastructure. For several good reasons, the Computer Society and other societies should control their own budgets, and central control should be minimal. This sense of ownership is essential for societies to carry out their missions of introducing new products and remaining vibrant.

Our technologies are becoming increasingly integrated. The IEEE should maintain the role of encouraging societies to work together to develop interdisciplinary member services.

As your president, I will support such interaction. I will work to eliminate central impediments, bring about society empowerment, and build strategic relationships among the societies. The real challenge here, of course, is to work for a common vision: to offer the best service to our members and customers. If we keep this in mind, the IEEE's societies and distributed management really are strengths.

Luis T. Gandía

During the Wall Street bull market, no one cared about these details. Returns on our investments were sufficient to take care of initiatives and cover part of our overhead. The market drop brought with it a new reality, and societies have had to bear the weight of the ensuing changes.

Societies are the backbone of the IEEE, and they produce our intellectual property. Therefore, they must have enough financial reserves to continue producing products and services for our members and for the society of professionals we serve.

Our bylaws state that all reserves belong to the IEEE. I am working on a proposal to change the IEEE bylaws to allow societies to retain part of the reserves they produce. This would empower societies to continue their work. While this will not improve the IEEE's financial situation, it would allow societies to invest part of their reserves in new products and services.

Arthur W. Winston

The Computer Society, the largest of the IEEE societies, has an extremely good volunteer and staff structure. Computer Society leadership is capable of determining and executing what they believe to be of most value to their members.

However, the same cannot be said for many of the other societies. IEEE leaders need to update the strategic plan, then determine what resources are needed to carry it out. They also must establish metrics and perform a cost/benefit evaluation of current spending trends.

Volunteers should assert greater control over what the IEEE should be doing and can afford. While the corporate IEEE has legal control of the organization's funds, the IEEE Board needs to delegate greater authority to the societies, allowing them to use their portion of the overall IEEE budget at spending levels commensurate with that delegation. If the $50,000 spending limit is too low, I would recommend that the Board raise it.

Question

2 Budgeting at the IEEE seems to be largely a matter of significantly increased corporate spending while generating too little revenue to cover that spending, and then balancing the budget through taxes on the societies and other successful IEEE units.

What would you do as president to manage the IEEE expense budget, including reducing corporate spending?

Computer Society Position

2 The Computer Society is very concerned that corporate costs are endangering the IEEE's viability. While Computer Society leaders agree that we should pay our fair share of corporate support expenses, we want reassurance that corporate expenses are reasonable and are managed efficiently.

IEEE societies generate the majority of IEEE intellectual property and revenue. Redirecting funds to cover central IEEE deficits threatens the societies' ability to serve members. The IEEE must enforce market discipline of balancing revenues with minimal overhead and bureaucracy.

The IEEE president should support the work of the Operations Review Committee, a group charged with reviewing corporate infrastructure costs and making recommendations for change. The president should direct the IEEE executive director to reduce corporate infrastructure charges and implement new financial practices, such as considering best-of-class costs when evaluating expenses.

The IEEE should eliminate operations best done by organizational units closest to the members.

Question

3 The IEEE is a global professional organization that serves a body of professionals and has historically been committed to the free and open communication of public scientific knowledge. Recent actions by the IEEE Intellectual Property Rights Office have eroded that commitment, and many researchers could decide to publish elsewhere rather than submit to IEEE copyright restrictions.

What specific steps would you take to ensure that the IEEE's traditional role in scientific communication continues?

3 The Society Board has passed a resolution urging the IEEE to investigate the applicability of export restrictions and to maximize the dissemination of open public scientific knowledge by all technological means. Through IEEE-USA, the IEEE should urge the US Congress and concerned government agencies to adopt more enlightened laws in this area. These policies should balance legitimate intellectual property rights and security concerns with the overarching need for the kind of open communication that sustains advances in scientific and engineering fields. These advances will in turn sustain democratic government.

Rather than change our operations to meet legal requirements, IEEE lawyers should advise the IEEE as to how to maintain our usual publishing operations as much as possible while still satisfying new laws.

Because the societies generate the majority of IEEE intellectual property, volunteer leaders from within the societies should be involved in any proposed decisions to change IEEE copyright policies.

Vijay K. Bhargava

2 IEEE budgeting is currently an annual exercise. As your president, I will first conduct a thorough and independent review of all internal operations, examining accounting controls, information systems charges, institutional research costs, travel services ledgers, legal fees, marketing costs, and operational expenses. I'll then plan a multiyear budget.

To reduce corporate spending, we need to: a) make our internal operations competitive with external providers, b) prioritize all activities and focus resources on those that are relevant, c) work with societies to simplify business rules, d) consider outsourcing when it is cost effective, e) have volunteer oversight on IT expenses, and f) approve major programs only after examining their value and long-term financial implications.

Implementing sound budgeting principles will ensure that the operations budget is balanced, new initiatives spending is limited to a portion of investment income, and direct and indirect cost allocation is equitable.

3 I assume that this question is driven by policies of the US Office of Foreign Assets Control and International Traffic in Arms Regulations. As a result of these regulations, the IEEE copyright form was drastically changed to become very restrictive. Now that the form has reverted to something very close to the original, our traditional role in scientific communication should resume. However, for our US members, the IEEE US Export Control Compliance Form still exists.

Our journals, conference proceedings, books, and short courses—our intellectual property—account for 75 percent of IEEE revenues. Protecting and increasing this income stream is a must. This may include turning the IEEE Web site into a portal for related technical material and brokering cooperative arrangements with other publishers. The IEEE Intellectual Property Rights Office needs to keep these objectives in mind and take steps that will encourage contributions from authors.

Luis T. Gandía

2 To develop a budget, it's necessary to plan ways to reduce costs, often by integrating similar operational tasks and eliminating redundancies. Ideally, the IEEE will be able to bring expenses in line with income and still preserve the quality of services. This becomes the joint responsibility of staff and volunteers. If a reduction in staff is required, it should be made based on well-defined objectives.

As president, I would recommend a review of our operations so that volunteers and staff can work more effectively to advance our programs within budgetary restrictions. As we invest in new programs, we need to consider the long-term implications of those investments.

Additionally, I would focus attention on considering the IEEE's priorities so that funding can be available for introducing new products and meeting the needs of our members. Our priorities need to take into account the value that would be gained from each investment of time and money.

3 This problem was created when staff, without consulting volunteers and with very conservative advice from legal counsel, changed the IEEE copyright form. As soon as our volunteer leaders involved with publications found out about this fiasco, they started working with staff and lawyers, convincing them that this was a great error. As a result, we are back to our old copyright form.

Indeed, this example clearly shows why volunteers must be kept abreast of important matters within the IEEE. Because our publications contain a wealth of intellectual property, we must make sure that our authors are happy to have their works appear in IEEE publications.

As president, I would look very closely at these kinds of actions and use my leadership position to make sure that key volunteers are informed and situations like this.

Arthur W. Winston

2 It is reasonable to question the current level of corporate expenditures. IEEE leaders should establish metrics to determine what program areas and expenses are most important for achieving the IEEE's strategic mission. Currently, the Board is not involved in reviewing either staffing levels or organizational structures that would best use our resources. Also, staff members often do not provide the kinds of timely and useful information that would enable the Board to make effective decisions.

As president, I will encourage the Board to become more actively involved in these matters. I will work with other volunteers to reduce expenses further by reviewing meeting venues, schedules, and support staff to bring spending in this area to a more reasonable level. We must approach this issue in a businesslike fashion, bearing in mind changes in the marketplace and in world dynamics.

3 At the June Transnational Committee meeting, we discussed the copyright issue, and the consensus was to oppose the restrictions. We were told that the restrictions would be removed. It is my understanding that most of the original IEEE copyright form has been restored.

The form has been split into two parts, so that the consideration of export control has been removed for most authors. However, IEEE leaders and US government representatives are working on the issue of export control for military or defense-related material. I believe that such security restrictions would apply to any organization or publication. But we should not let these restrictions creep in where they are not necessary.

Staff apparently acted without sensitivity to volunteer needs in this case. This incident shows the danger of being led by staff without volunteer oversight.

Question

4 A fundamental tenet of IEEE governance historically has been that the IEEE is a volunteer-led and volunteer-driven organization. However, since 2000, IEEE policies and practices seem to be directed more by legal counsel and corporate staff than by volunteer leaders.

In your view, what should be the role of volunteers in the management of the IEEE, and what will you do as president to reinforce that role? How will you address concerns that the balance of decision-making authority in the IEEE has shifted away from its volunteer leaders and reverse the trend of increasing corporate staff control?

Computer Society Position

4 A strong professional staff operation is as essential to the IEEE's success as it is to the Computer Society's well-being. However, as a volunteer-driven organization, the president and the Board should determine IEEE policies and overall strategic directions. IEEE corporate staff members should support volunteer leaders and implement volunteer policies.

The IEEE president should clarify and reinforce the roles of volunteer leaders and their staff counterparts to avoid conflicts in strategic direction. The president should seek the counsel of staff and outside legal consultants, but understand that these are recommendations to consider rather than absolutes to follow. Further, the president should use such recommendations to manage operations so that they meet legal requirements, rather than overhaul operations to comply with rigid legal interpretations. The president should not delegate to corporate staff or legal counsel his responsibility to develop policies in the best interests of IEEE members.

Question

5 IEEE societies must have extensive volunteer participation in order to meet the needs of their members. A cornerstone of the IEEE economic model is the presence of a strong society volunteer base. Simply put, the IEEE and its societies can neither maintain current products and services nor embrace new ones without the dedication and commitment of a large group of active volunteers. Recruiting and retaining volunteers and volunteer leaders for the societies is becoming increasingly difficult.

As president, how will you reinforce a strong volunteer base in societies? What will you do to support society volunteers in their important work and help make their jobs easier?

5 The Computer Society's broad array of programs can only be sustained through the efforts of many volunteers. Its volunteer workforce is "paid" in the personal satisfaction derived from contributions to the Society and the profession. But, because of increased IEEE corporate centralization, volunteer satisfaction is becoming harder to achieve. Indeed, the Computer Society has already lost senior volunteer leaders who could not accept unnecessary changes that the IEEE dictated. IEEE leaders imposed those changes without consulting societies and apparently without carefully thinking through the consequences.

The IEEE president should direct leaders and staff to improve areas where the lack of authority or failed communication has impaired volunteer participation. The president should require IEEE vice presidents and the executive director to consult with society volunteers before enacting policies that can affect societies. The president should reinforce the principle that policies represent prudent controls rather than burdensome requirements.

Vijay K. Bhargava

4 I will work with society leaders to counter the legal and accounting obstacles that are getting in the way of effective volunteer leadership. The IEEE works best when there is excellent synergy between volunteers and staff and across entities. Communication is the key.

We must remain a volunteer-led and volunteer-driven organization. I will reinforce this concept by ensuring that our staff understands it. I would also ensure that our volunteers understand that we can count on staff to carry out IEEE policies but that we should not burden them with our occasional near-term, parochial, or tactical requests. We need to work as a team to understand the challenges facing the IEEE. We need to work together strategically, positioning the IEEE for the future so that volunteer policies and staff implementation result in better products and services for our members.

5 This is currently one of the topics that the IEEE Strategic Planning Committee is addressing. There is a growing concern that the IEEE is becoming an "old boys/girls club." We need to involve more people in IEEE activities. The people are out there; the key is to identify them. Society chapters could play a strong role in identifying new volunteers.

I would also recommend that society Nominations Committees have at least one representative from Regions 8, 9, and 10 (that is, countries outside North America). This will allow us to draw from a vast and active member base in these regions.

We need to expand membership beyond the traditional field of electrical engineers. There are tremendous opportunities in information technology and related areas. This too should help the societies strengthen their volunteer base. I will support these volunteers by working to simplify bylaws and delegating authority to counter staff control.

Luis T. Gandía

4 Without a doubt, our staff leadership seeks legal advice more now than in the past. While previous executive directors conducted business without legal counsel, our current executive director believes in seeking legal advice before making any major decision. This practice increases the IEEE's expenses, often excludes volunteer input, and could transfer important decisions out of our hands.

In my view, the IEEE Board of Directors should determine the path our IEEE is to follow. The duty of staff members is to carry out instructions from the Board, and our volunteers must demonstrate leadership by making sure that this happens. Volunteers who have management experience are well suited to providing this type of oversight. Open and improved communication between key volunteers and staff could help make this possible.

5 Indeed, recruiting and retaining volunteers and leaders is becoming difficult. Lifestyles have changed dramatically, and employees are often obligated to work longer hours.

The IEEE has not been particularly successful in addressing this recurring topic. In order to attract more volunteers and leaders, we must put forth a major effort, articulating our objectives and obtaining adequate funding. Our progress so far has been proportional to our investment, and that investment has been minimal. There have only been sporadic efforts, largely driven by organizational units.

We have not sold academia or industry on the idea that active participation in the IEEE is beneficial to their employees. Our proposed means of recruitment lack compelling proof that would convince an organization that the investment is justified. We need to demonstrate, in very specific terms, how our offerings benefit academia or industry.

Arthur W. Winston

4 I am on record as having voiced concerns that volunteer control over the organization's direction has diminished.

I witnessed this trend toward being directed by legal counsel and corporate staff even before 2000. In the late 1990s, IEEE lawyers effectively hampered steps that would have allowed more responsibility and authority to be passed down to the organizational units. Similarly, as a past chair of the Audit Committee, I witnessed attempts at staff-led financial control.

During my days in business, I first determined what I wanted to accomplish and then brought in accountants and lawyers to implement my plan and ensure that my company's actions were legal. In the business world, I learned not to allow these consultants to lead the process or make business decisions for my company. As IEEE president, I would use the same approach. For decisions that involve risk, I would ask legal and accounting advisers for a risk assessment.

5 The societies are the backbone and the volunteers are the muscles of the IEEE. As president, I will work with the Technical Activities Board and society officers to keep the societies financially viable by controlling corporate charges and minimizing excessive corporate control. As in all my activities, I would make myself available to assist the societies in outreach activities.

Draft Computer Engineering Curriculum Volume Available for Review

The IEEE Computer Society, together with the ACM, has had a long history of involvement in preparing curriculum guidelines and tools for computer engineering academic programs. Updated gradually over many years, the current curriculum project, Computing Curricula 2001, collects resources from a wide range of computer science and engineering disciplines to present model courses, curricula, and body of knowledge documents for use in college programs. A substantial volume from Computing Curricula 2001 is now available for public review.

The preliminary draft of *Computing Curricula 2001: Computer Engineering*, available online at http://www.eng.auburn.edu/ece/CCCE/, is open for comments from the professional community. Academics, researchers, and representatives from industry are welcome to review the document and provide feedback to help shape the final volume.

Computer engineering is one of four content areas covered by the curriculum project. Other planned curriculum documents include volumes on computer science, software engineering, and information systems.

In addition to the review version of computer engineering, the *Computing Curricula 2001: Computer Science* volume is available for purchase at http://computer.org/cspress/CATALOG/cs01499.htm.

Edsger Wybe Dijkstra, 1930-2002

Edsger Wybe Dijkstra, a computer pioneer and emeritus professor at the University of Texas, died in August in Nuenen, the Netherlands. Trained in mathematics and physics as well as a computer science, Dijkstra was perhaps best known for his insistence that computer programs should be based on mathematical logic.

In the late 1950s, Dijkstra contributed to the development of Algol. For its clarity of purpose and its mathematical base, many look to this as a model for excellence in high-level programming languages.

Noted for the simplicity and elegance of his programming, Dijkstra also coined many phrases now common to computing parlance, including "structured programming." As noted in the *Oxford English Dictionary*, he introduced the terms "vector" and "stack" to their use in a computing context.

Throughout his professional life, Dijkstra circulated technical notes, trip reports, and other observations and commentaries—known to recipients as *EWDs*—to an audience that eventually numbered in the thousands. These works are now archived online at http://www.cs.utexas.edu/users/EWD/.

Dijkstra's awards and honors include the IEEE Computer Society's Harry M. Goode Memorial Award and Pioneer Award, the ACM Turing Award, and the distinction of Fellow of the Netherlands Royal Academy of Arts and Sciences, the British Computer Society, and the ACM. He was named a Foreign Honorary Member of the American Academy of Arts and Sciences in 1975.

Kelly Named IEEE Computer Society Associate Executive Director

Anne Marie Kelly, longtime member of the IEEE Computer Society staff, was recently selected as associate executive director. In addition to maintaining her role as the director of volunteer services, Kelly will represent and act in the place of Executive Director David W. Hennage when he is not available.

Kelly joined the Computer Society in 1988 as director of conferences. In 1994, this department was expanded to include staff support for the Society's activities in standards, technical committees, chapters, and student and educational programs. Kelly has headed the renamed Volunteer Services Department since that time. However, she stepped aside from those duties for nearly two years to serve the Society as acting executive director, from January 2000 until Hennage's appointment in November 2001.

In 1998, the Society Board of Governors presented Kelly with the Harry Hayman Award for Distinguished Staff Achievement. One of only five Hayman Award winners, Kelly was awarded a second Hayman plaque in 2001 for her outstanding performance as acting executive director.

Editor: Mary-Louise G. Piner, *Computer*, 10662 Los Vaqueros Circle, PO Box 3014, Los Alamitos, CA 90720-1314; mpiner@computer.org

CALLS FOR IEEE CS PUBLICATIONS

The IEEE Computer Society is seeking articles for *IEEE Security & Privacy,* a new magazine to be launched in January 2003. The primary objective of *IEEE Security & Privacy* is to stimulate and track advances in information assurance and security and present these advances in a form that can be useful to professionals ranging from academic researchers to industry practitioners.

IEEE Security & Privacy is envisioned to cover diverse aspects of information assurance such as legal and ethical issues, privacy concerns, tools to help secure information, analysis of vulnerabilities and attacks, trends and new developments, pedagogical and curricular issues in educating the next generation of security professionals, secure operating systems and applications, security issues in wireless networks, design and test strategies for secure and survivable systems, and cryptology.

For more information, see http://computer.org/computer/sp/.

IEEE Pervasive Computing seeks papers focusing on hardware technology, software infrastructure, real-world sensing and interaction, human-computer interaction, or systems considerations. Papers should range from 4,000 to 6,000 words (tables and figures count as 250 words each). For more information, see http://computer.org/pervasive/author.htm or send e-mail to pervasive@computer.org.

IEEE Internet Computing seeks original articles for a special issue on developments in IPv6 protocols and applications. Topics of interest include operating system support for IPv6; comparing IPv4-to-IPv6 transition mechanisms; routing protocol implementations; mobile IPv6 protocols and services; service provider, enterprise, and home networking environments and applications; migrating IPv4 applications to IPv6; implementing IPv6 in emerging 3G wireless networks; IPv6 anycast addressing for autoprovisioning and discovery; operational experiences on IPv6 networks; QoS and flow-label semantics; IPv6 packet header processing in routers; and IPv6 and MPLS.

The submission deadline is **1 Oct.** Contact guest editor Chris Metz, chmetz@cisco.com, and see http://computer.org/internet/edguide.htm for more information.

IEEE Computer Graphics and Applications seeks submissions for a July/August 2003 issue on non-photorealistic rendering (NPR). While computer graphics has traditionally centered on the production of photorealistic images to depict 3D scenes, NPR investigates alternatives that use abstraction and stylization to evoke the complexity of such scenes indirectly.

Areas of interest include simulation of natural media, traditional illustration effects, technical or medical illustration, real-time NPR algorithms, and animation and temporal coherence.

The deadline is **25 October.** Direct correspondence to guest editors Adam Finkelstein and Lee Markosian at cga-

Submission Instructions

The Call and Calendar section lists conferences, symposia, and workshops that the IEEE Computer Society sponsors or cooperates in presenting. Complete instructions for submitting conference or call listings are available at http://computer.org/conferences/submission.htm.

A more complete listing of upcoming computer-related conferences is available at http://computer.org/conferences/.

Call for Articles for *Computer*

Computer seeks articles for a special issue on hardware-software codesign, to appear in April 2003. Guest editors are Jörg Henkel, NEC; Xiaobo Sharon Hu, University of Notre Dame; and Shuvra S. Bhattacharyya, University of Maryland.

Through a decade of intense research, hardware-software codesign has become an integral part in the design process of embedded systems. The advent of hardware-software codesign techniques has facilitated shorter times to market, reduced system costs, and increased functionality and complexity for many of today's embedded systems designs. However, as the demand for more complex, faster, lighter, and cooler embedded systems is ever increasing, the codesign field is facing even bigger challenges.

The goal of this special issue is to publish cutting-edge research in hardware-software codesign techniques, tools, and architectures that is aimed at providing novel solutions for the design process of complex embedded systems such as mobile devices. Proposed topics include

- computer-aided codesign techniques: specification and modeling, synthesis, partitioning, estimation, platform-based design, design space exploration, co-simulation, and test strategies;
- software for codesign: code generation, software development environments, RTOS, process scheduling, system integration, communication protocols; and
- codesign architectures: hardware-software interfaces, distributed and multiprocessor architectures, and reconfigurable architectures.

The submission deadline is **15 October.** Author guidelines are available at http://computer.org/computer/author.htm. Send an electronic file to Jörg Henkel, http://www.codesign-symposium.org. Send queries to henkel@nec-lab.com.

npr@cs.princeton.edu; send manuscripts to cga-ma@computer.org. See http://www.npar.org/cga/.

The July-September 2003 issue of *IEEE Annals of the History of Computing* will be devoted to recording and recounting efforts to preserve computing practice through physical reconstruction, restoration, and simulation of historic computing devices. In addition to providing accounts of such projects through substantive articles, the issue is intended to serve as a digest or register of projects and initiatives, past and present, as well as societies and organizations active in these areas. Doron Swade (d.swade@nmsi.ac.uk) will serve as guest editor for this special issue on historical reconstructions. Submissions are due by **30 Nov.** For the full call, see http://computer.org/annals/call.htm.

Communication between human users and their computer systems has always been a crucial component of interactive graphics applications. The September/October 2003 issue of *IEEE Computer Graphics and Applications* will focus on recent advances in perceptual multimodal interfaces. Potential topics include design approaches for multimodal interaction systems, human performance and perception studies devoted to multimodal technologies, user awareness and context-based response at the interface, and performance metrics and evaluation studies.

Submissions are due to cga-ma@computer.org by **1 December.** Direct correspondence to both guest editors: L. Miguel Encarnação, me@crcg.edu, and Lawrence J. Hettinger, lhettinger@northropgrumman.com. See http://computer.org/cga/CFPSep03.htm.

Massive amounts of 3D data are now available. To make these data useful, they should represent the real world in models available for interactive exploration and analysis. Problems related to fusing multisource data make this modeling problematic. The November/December 2003 special issue of *IEEE Computer Graphics and Applications* will feature 3D reconstruction and visualization of large-scale environments.

Potential topics include issues of scale, problems of dynamic data, updating models with newly acquired data, and incomplete models.

Submit papers to cga-ma@computer.org by **1 March 2003.** Direct correspondence to guest editors William Ribarsky, ribarsky@cc.gatech.edu, and Holly Rushmeier, hertjwr@us.ibm.com. See http://computer.org/cga/CFPNov03.htm.

IEEE Annals of the History of Computing Editor in Chief Search

The IEEE Computer Society seeks, by **15 September,** nominations and applications for the volunteer position of editor in chief of *IEEE Annals of the History of Computing*.

From the analytical engine to the supercomputer, from Pascal to von Neumann, from punched cards to CD-ROMs—the *IEEE Annals of the History of Computing* covers the breadth of computer history.

The EIC is appointed by the IEEE Computer Society president for a two-year term, and the position can be extended for a second term. The EIC is responsible for the volunteer leadership of *Annals* including soliciting high-quality manuscripts, coordinating and overseeing the peer-review process, selecting and assisting guest editors of special issues, recommending candidates for the editorial board, and chairing the editorial board. For more information, contact search committee chair Peggy Aldrich Kidwell at kidwellp@si.edu or see http://computer.org/annals/an2002/CfEIC.lo.pdf.

Search for Editor in Chief of the IEEE Computer Society Web Site

The IEEE Computer Society is seeking applicants, by **11 October**, for the position of editor in chief of the IEEE Computer Society Web site. The EIC will serve a two-year term, renewable for a second two years.

The Computer Society Web site, http://computer.org, presents both informative pages on standing Computer Society programs and dynamic features such as the new distance learning campus of free courses for members. Consisting of thousands of subpages that discuss our standards activities, technical forums, student scholarship programs, and conference dates, computer.org also contains a digital library of 19 magazines and transactions, with issues available dating back to 1988.

The major responsibilities of the Web site EIC include reviewing areas of the Web site, participating in discussions about the Society's online future, and supervising an upcoming site redesign.

If you have any questions, suggestions of possible candidates, or would like to indicate interest, please contact Jim Aylor, Chair, EIC Search Committee, at jha@virginia.edu.

OTHER CALLS

CSEET 2003, 16th Int'l Conf. on Software Eng. Education & Training, 20-22 Mar. 2003, Madrid, Spain. Submission deadline **1 Oct.**; see http://www.ls.fi.upm.es/cseet03/.

PERCOM 2003, Int'l Conf. on Pervasive Computing & Comm., 23-26 Mar. 2003, Ft. Worth, Texas. Submission deadline **1 Oct.** Contact Mohan Kumar, kumar@cse.uta.edu. See http://percom.org/.

ICSE, Int'l Conf. on Software Eng., 3-10 May 2003, Portland, Ore. Submission deadline **4 Oct.;** see http://www.icse-conferences.org/2003/.

IPDPS 2003, 17th Int'l Parallel & Distributed Processing Symp., 22-26 Apr. 2003, Nice, France. Papers due **4 Oct.** See http://www.ipdps.org/ipdps2003/2003_cfp.htm.

CSMR 2003, 7th European Conf. on Software Maintenance and Reengineering, 26-28 Mar. 2003, Benevento, Italy. Papers due **10 Oct.** See http://rcost.unisannio.it/csmr2003.

ARITH-16, 16th IEEE Symp. on Computer Arithmetic, 15-18 June 2003, Santiago de Compostela, Spain. Electronic submission deadline is **15 Oct.** Contact Jean-Claude Bajard, bajard@lirmm.fr. See http://www.dec.usc.es/arith16/.

CCGrid 2003, IEEE Int'l Symp. on Cluster Computing & the Grid, 11-16 May 2003, Tokyo. Papers due **30 Oct.** See http://www.ccgrid.org/.

ISMVL 2003, 33rd Int'l Symp. on Multiple-Valued Logic, 16-19 May 2003, Tokyo. Submissions due **1 Nov.** See http://science.icu.ac.jp/ismvl/.

CVPR, Computer Vision & Pattern Recognition Conf., 16-22 June 2003, Madison, Wis. Papers due **4 Nov.** See http://www.cs.toronto.edu/cvpr2003/.

Metrics 2003, 9th Int'l Software Metrics Symp., 3-5 Sept. 2003, Sydney, Australia. Abstracts due **15 Dec.** See http://metric.cse.unsw.edu.au/Metrics2003/documents/Metrics2003cfp.pdf.

FIE 2003, Frontiers in Education Conf., 5-8 Nov. 2003, Boulder, Colo. Abstracts due **6 Jan. 2003.** See http://fie.engrng.pitt.edu/.

MSE 2003, Int'l Conf. on Microelectronic Systems Education, 1-2 June 2003, Anaheim, Calif. Papers due **15 Jan. 2003.** Contact Don Bouldin, dbouldin@tennessee.edu. See http://www.mseconference.org/mse03-cfp.pdf.

CALENDAR

OCTOBER 2002

1 Oct: SCAM 2002, 2nd Int'l Workshop on Source Code Analysis & Manipulation, (with ICSM), Montreal. See http://www.brunel.ac.uk/~csstmmh2/scam2002/.

1 Oct: DBMR 2002, Database Maintenance and Reengineering Workshop, (with ICSM), Montreal. See http://alarcos.inf-cr.uclm.es/dbmr2002/principaldmr.htm.

2 Oct: WSE 2002, Workshop on Web Site Evolution, (with ICSM), Montreal. See http://star.itc.it/wse2002/.

2 Oct: WESS 2002, 8th IEEE Workshop on Empirical Studies of Software Maintenance, (with ICSM), Montreal. See http://www.info.uqam.ca/~lounis/wess2002_cfp.htm.

3-4 Oct: ISESE 2002, Int'l Symp. on Empirical Software Eng., Nara, Japan. See http://se.aist-nara.ac.jp/isese2002/.

3-6 Oct: ICSM 2002, IEEE Int'l Conf. on Software Maintenance, Montreal. See http://www.icsm2002.org.

7-10 Oct: ISWC 2002, 6th Int'l Symp. on Wearable Computers, Seattle, Wash. See http://iswc.tinmith.net/.

8-10 Oct: ITC, 2002 Int'l Test Conf., Baltimore. See http://www.itctestweek.org/.

9-12 Oct: Ubiquity, Grace Hopper Celebration of Women in Computing Conf., Vancouver, Canada. Contact taylor@ece.northwestern.edu. See http://www.gracehopper.org/.

11-13 Oct: DS-RT 2002, 6th IEEE Int'l Workshop on Distributed Simulation & Real-Time Applications, Ft. Worth, Texas. See http://www.cs.unt.edu/~boukerch/DS-RT2002.

12-16 Oct: IEEE MASCOTS 2002, 10th IEEE Symp. on Modeling, Analysis, & Simulation of Computer

Search for Editor in Chief of *IEEE Security & Privacy*

The IEEE Computer Society is seeking applicants, by **1 October,** for the position of editor in chief of *IEEE Security & Privacy,* a new magazine to be launched in January 2003. The first EIC will serve a two-year term, renewable for a second two years.

IEEE Security & Privacy magazine is being planned to address a broad range of topics related to securing information and computing resources. It is intended to serve as one of the premier channels for dissemination of results from leading-edge research, practice, and experience in various aspects of information assurance.

The major responsibilities of the first EIC of *IEEE Security & Privacy* include soliciting manuscripts from potential authors, guiding the review process to ensure the publication of high-quality work, identifying potential editorial board members, and selecting competent reviewers.

If you have any questions or suggestions for possible candidates, or if you would like to indicate interest in the position, please contact James Davis, Chair, EIC Search Committee, at davis@iastate.edu.

For further details on the search for editor in chief, see http://computer.org/pr/Aug02/SP_EIC.htm.

& Telecommunications Systems, Ft. Worth, Texas. See http://www.cs.unt.edu/~boukerch/mascots2002.

13-16 Oct: SRDS 2002, 21st Symp. on Reliable Distributed Systems, Suita, Japan. See http://www-kiku.ics.es.osaka-u.ac.jp/SRDS2002.html.

14-16 Oct: ICMI 2002, Int'l Conf. on Multimodal Interfaces, Pittsburgh. See http://www.is.cs.cmu.edu/icmi.

16-18 Oct: AIPR 2002, The 31st Applied Imagery Pattern Recognition Workshop, Washington, D.C. See http://www.aipr-workshop.org.

23-25 Oct: HASE 2002, 7th IEEE/IEICE Int'l Symp. on High-Assurance Systems Eng., Tokyo. See http://unicorn.eei.metro-u.ac.jp/hase2002/.

24-25 Oct: MA 2002, 6th Int'l Conf. on Mobile Agents, Barcelona, Spain. See http://dmag.upf.es/mata-ma2002/.

27-29 Oct: HLDVT 2002, High-Level Design Validation & Test Workshop, Cannes, France. See http://tima.imag.fr/conferences/hldvt/index.html.

27 Oct.-1 Nov: VIS 2002, IEEE Visualization Conf., Boston. See http://vis.computer.org/vis2002.

28-29 Oct: InfoVis 2002, IEEE Symp. on Information Visualization, (with VIS 2002), Boston. See http://www.infovis.org/infovis2002/.

28-29 Oct: VolVis 2002, IEEE/SIGGRAPH Symp. on Volume Visualization & Graphics, (with VIS 2002), Boston. See http://www.cs.sunysb.edu/~volvis02/.

Online Calendar and Calls

For more information on calls for papers, conferences, workshops, and symposia, check our full listing online at http://computer.org/conferences/.

28-30 Oct: FAABS II, 2nd Goddard Workshop on Formal Approaches to Agent-Based Systems, Greenbelt, Md. See http://fmw.gsfc.nasa.gov.

NOVEMBER 2002

4-6 Nov: ICTAI 2002, 14th IEEE Int'l Conf. on Tools with AI, Manassas, Va. See http://www.nvc.cs.vt.edu/ictai2002/.

6-7 Nov: IWADS 2002, Int'l Workshop on Autonomous Decentralized Systems, Beijing, China. See http://iwads02.njtu.edu.cn/.

6-8 Nov: 17th IEEE Int'l Symp. on Defect & Fault Tolerance in VLSI Systems, Vancouver, Canada. See http://www.elet.polimi.it/dft/.

6-8 Nov: Cyber Worlds: Theories & Practices Int'l Symp., Tokyo. See http://cis.k.hosei.ac.jp/CW2002/.

6-8 Nov: LCN 2002, 27th IEEE Conf. on Local Computer Networks, Tampa, Fla. See http://www.ieeelcn.org.

6-9 Nov: FIE, Frontiers in Education, Boston. See http://www.wpi.edu/News/Conf/FIE2002/.

10-14 Nov: ICCAD 2002, Int'l Conf. on Computer Aided Design, San Jose, Calif. See http://www.iccad.com.

12-15 Nov: ICNP 2002, 10th IEEE Int'l Conf. on Network Protocols, Paris. See http://protocols.netlab.uky.edu/icnp/index.html.

12-15 Nov: ISSRE 2002, 13th Int'l Symp. on Software Reliability Eng., Annapolis, Md. See http://www.issre2002.org.

16-19 Nov: FOCS 2002, 43rd Ann. IEEE Symp. on Foundations of Computer Science, Vancouver, Canada. See http://www.mitacs.math.ca/focs02/.

17-22 Nov: SC 2002, Baltimore. See http://www.sc-2002.org/.

18-20 Nov: ATS 2002, 11th IEEE Asian Test Symp., Tamuning, Guam. See http://ats02.ip.elec.mie-u.ac.jp/.

18-22 Nov: MICRO, 35th Ann. Int'l Symp. on Microarchitecture, Istanbul. See http://www.microarch.org/micro35/.

21-22 Nov: WRTLT 2002, 3rd Workshop on RTL ATPG & DFT, (with ATS), Tamuning, Guam. See http://www.ip.elec.mie-u.ac.jp/~wrtlt02/.

22 Nov: WWC 2002, 5th Ann. Workshop on Workload Characterization, Austin, Texas. See http://www.ece.utexas.edu/~ljohn/wwc.html.

DECEMBER 2002

3-4 Dec: 6th IEEE Workshop on Applications of Computer Vision, Orlando, Fla. See http://www.cs.ucf.edu/~vision/workshop/2002/applicationCompVision.html.

3-5 Dec: RTSS 2002, 23rd IEEE Int'l Real-Time Systems Symp., Austin, Texas. See http://www.cs.unc.edu/rtss2002.

4-6 Dec: 27th Ann. IEEE/NASA Software Eng. Workshop, Greenbelt, Md. See http://sel.gsfc.nasa.gov.

5-6 Dec: Motion 2002, Workshop on Motion & Video Computing, Orlando, Fla. See http://www.cs.ucf.edu/~vision/workshop/2002/motionVideoComp.html.

8 Dec: WIESS 2002, 2nd Workshop on Industrial Experience with Systems Software, Boston. See http://www.usenix.org/events/wiess02/.

9-11 Dec: OSDI 2002, 5th Symp. on Operating Systems Design & Implementation, (collocated with WIESS 2002), Boston. See http://www.usenix.org/events/osdi02/.

Risk-Based Approach to Mission-Critical Software Testing

Systematic Software Testing, Rick D. Craig and Stefan P. Jaskiel. This book offers those involved in building and maintaining complex, mission-critical software systems a flexible, risk-based process to improve their software testing capabilities.

The authors describe how to use a preventive testing method that parallels the software development life cycle, then explain how to create and subsequently use test plans, test design, and test metrics. They present detailed instructions to help readers decide what to test, how to prioritize tests, and how to determine when testing is complete. The book also covers risk analysis and measuring test effectiveness to maximize testing efficiency.

Artech House, Boston; http://www.artechhouse.com; ISBN 1-58053-508-9; 568 pp.; $69.

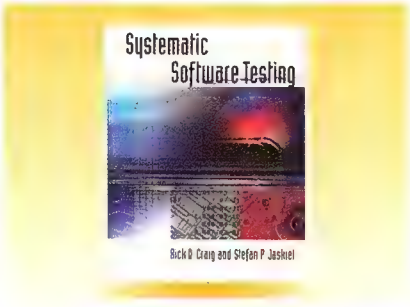

TLA+ SPECIFICATION LANGUAGE PRIMER

Specifying Systems: The TLA+ Language and Tools for Hardware and Software Engineers, Leslie Lamport. This book shows how to write unambiguous specifications of complex computer systems. The author introduces specification by explaining with mathematical precision how to describe a system's behavioral properties, emphasizing safety properties. Lamport covers more advanced topics, including liveness and fairness, real-time properties, and composition.

The last half of the book provides a complete reference manual and mini-manual for the TLA+ language. Tools for TLA+ syntax analysis and model checking can be found on the Web, along with supplemental materials and exercises.

Addison Wesley, Boston; http://www.awl.com/cseng; ISBN 0-32114-306-X; 384 pp.; $44.99.

INTERPOLATION TUTORIAL

Nonuniform Sampling: Theory and Practice, Farokh Marvasti, ed. This book's 20 chapters contain contributions by leading researchers in nonuniform and Shannon sampling, zero crossing, and interpolation theory. Its practical applications include seismology, speech and image coding, modulation and coding, optimal content, array processing, and digital filter design.

The book's tutorial approach targets practicing engineers and advanced students in science, engineering, and mathematics. It can also serve as a reference for scientists and engineers working in the areas of medical imaging, geophysics, astronomy, biomedical engineering, computer graphics, digital filter design, speech and video processing, and phased array radar. The accompanying CD-ROM contains C-codes and Matlab and Mathcad programs for the book's algorithms.

Kluwer Academic/Plenum Publishers, New York; http://www.wkap.nl/; ISBN 0-306-46445-4; 912 pp.; $145.

PATTERN-MATCHING TOOLKIT

Flexible Pattern Matching in Strings: Practical On-Line Search Algorithms for Texts and Biological Sequences, Gonzalo Navarro and Mathieu Raffinot. String matching problems range from searching a single text for a string of characters to searching a database for approximate occurrences of a complex pattern. Recently, interest in sophisticated string matching problems has increased dramatically, especially in information retrieval and computational biology. The authors present a practical approach to string matching problems, focusing on the algorithms and implementations that perform best in practice.

The book also covers searches for simple, multiple, and extended strings and for regular expressions. Detailed explanations, step-by-step examples, algorithm pseudocode, and implementation efficiency maps give readers the tools to choose appropriate algorithms for their applications.

Cambridge University Press, Cambridge, England; http://www.cambridge.org; ISBN 0-52181-307-7; 232 pp.; $50.

Editor: Michael J. Lutz, Rochester Institute of Technology, Rochester, NY; mikelutz@mail.rit.edu. Send press releases and new books to *Computer*, 10662 Los Vaqueros Circle, Los Alamitos, CA 90720; fax +1 714 821 4010; newbooks@computer.org.

Enhanced Bluetooth Test Tool from Tektronix

BPA105 is the latest version of Tektronix's Bluetooth Protocol Analyzer. The company began offering the BPA100 series soon after the Bluetooth specification became available, and it continues to enhance the product as the standard evolves. With Bluetooth product compliance gaining significant momentum, engineers face new design challenges as more Bluetooth-compliant devices enter the market. The new tool analyzes product interaction within a Bluetooth piconet, making it easier to identify and solve interoperability issues.

BPA105 can assist throughout the entire development cycle, including hardware and software design, debug, integration, precertification, and conformance testing. You can use the tool to decode BNEP, AT, and HID packets as well as log, decode, and display all base-band packets transmitted and received within a Bluetooth piconet. In piconet mode, you can introduce user-generated traffic, such as data packets, into transmissions.

The upgrade also features a sync wizard, audio support, and search and clock-out functionality, and you can cross-trigger the product with other test and debug equipment, such as a logic or spectrum analyzer.

Contact the company for licensing information; http://www.tektronix.com.

NEXIQ Technologies' IVIS-Lite lets you design, prototype, test, and produce information displays with limited memory and processor resources.

NEXIQ Technologies' Intelligent Display Software

IVIS-Lite from NEXIQ Technologies is a software package for informational display applications designed specifically for hardware platforms with limited memory and processor resources. A C-language interface facilitates integration of existing code, compilers, and development environments. You can use the software package to design, prototype, test, and produce displays for automobiles, industrial and manufacturing equipment, consumer electronics, home appliances, and heavy vehicles such as construction and agricultural equipment.

The product includes In-Vehicle Information System core functionality, which enables telematics services in motor vehicles and is traditionally used for applications with high-resolution displays and significant multitasking and data-storage requirements.

Designed for smaller, multiscreen information displays, IVIS-Lite performs on 8- and 16-bit processors and requires only 256 Kbytes of memory. The software features real-time antialiasing and supports color or monochrome displays of up to VGA resolution. Operating systems are supported on target hardware platforms but not required.

Contact the company for licensing information; http://www.nexiq.com/.

Actel's FPGA Development IDE

Actel has upgraded its IDE for field-programmable gate array development and design. Libero 2.2 features an incremental timing analysis engine and automated register retiming, which produce optimized circuits to make timing estimations more accurate. The software helps eliminate the labor-intensive process of analyzing critical paths and changing HDL code, and it can automatically reposition registers within combinatorial logic to balance routing and improve circuit performance.

The product now integrates with SynaptiCAD's WaveFormer, which lets you describe a simulation's stimulus graphically and then convert this data into a VHDL or Verilog test bench. WaveFormer includes support for VHDL generic statements as well as VHDL attributes for specifying port size. In Verilog, you define clocks in separate modules, which can help reduce the generated code's size and improve readability. The tool uses a btim binary waveform file format to handle larger files.

Actel Libero 2.2 is available in four versions. Libero Platinum costs $2,495 and offers unlimited design capacity and customer support, while Libero Gold costs $995 and is aimed at users designing system-level devices of 50,000 gates or fewer. Libero Silver and an evaluation version are available free of charge for one year and 45 days, respectively; http://www.actel.com.

Parasoft's Automated Java Classes Testing Unit

Jtest from Parasoft, a development tool that automates all aspects of unit testing for Java applications including white box, black box, and regression testing, now integrates with Sun Microsystems' Sun ONE Studio IDE.

In the rush to get a product to market, developers often overlook the tedious, time-consuming process of manually writing test cases for each Java class. In addition to automatically enforcing more than 300 coding standards through static analysis, the product has rules support for Enterprise JavaBeans components, Design by Contract, JavaServer pages technology, servlets, and project metrics.

Jtest costs $3,495 and is available for Solaris, Windows, and Linux; free evaluations can be downloaded from the company site at http://www.parasoft.com.

Packeteer Upgrades Central Reporting Application

ReportCenter 2.0 from Packeteer is a central reporting application designed to work with the company's PacketShaper and AppVantage systems. The application can generate detailed views

on more than 40 network and performance metrics that cover availability, throughput, use, efficiency, response-time, and service-level exceptions on a WAN, LAN, and application-class basis.

The product's data management capabilities range from fine-grained, up-to-the-minute reports to historical and ad hoc analysis and reporting that can sift through several years of data to provide insight into specific performance issues. You can view all reports through a centralized Web portal interface designed to facilitate navigation and drill down.

ReportCenter 2.0 runs on Windows NT 4.0, Windows 2000, and Windows XP. Prices vary by size of deployment. Contact the company for licensing information; http://www.packeteer.com.

VisiComp Releases Java Debugger

RetroVue from VisiComp is a Java debugger that lets you rewind or fast-forward through a program's execution to search for particular conditions or variables and observe how the source code, state, and threads interact.

The debugger works with long execution journals and displays results in a format that allows easy pattern recognition, reducing the time required to trace specific bugs back to their root cause and giving managers the freedom to assign the best developer to each problem. In addition to a project save-load capability, the product includes a GUI as well as a command-line launcher.

RetroVue costs $995; http://www.visicomp.com.

Compuware's Driver Development Suite for Windows

DriverStudio 2.7 from Compuware is designed to accelerate the development and deployment of Windows device drivers that meet strict WHQL (Windows Hardware Quality Labs) standards and tight time-to-market requirements.

The integrated tool suite includes SoftICE, DriverWorks, VtoolsD, a DriverNetworks frameworks package, and driver editions of BoundsChecker, TrueTime, and TrueCoverage. The upgrade introduces a host-target architecture that lets you rapidly debug, test, and tune drivers on a remote machine via a serial cable or TCP/IP connection. Other new features include full support for Windows XP, SoftICE support for Intel's hyper-threaded processors, and enhanced memory and resource tracking.

DriverStudio 2.7 costs $2,499; http://www.compuware.com.

Silas Technologies Upgrades Application-Monitoring Software

Silas Technologies has ugraded its application-monitoring software to provide access to real-time as well as historical data about the operational health of applications and affiliated resources across legacy mainframe, Unix, Windows, and Web servers. Silas Reveille 2.0 analyzes and reports on service-level agreements, real-time application and component status, resource specifics, and Web site performance.

The upgrade lets you conduct trending analysis on applications and their underlying components. A new management-reporting distribution feature delivers e-mail that summarizes application and resource-level activities. The product also includes additional database support and wizards for generating tests for applications, technology services, and other infrastructure components.

Silas Reveille 2.0 pricing ranges from $65,000 to $300,000; http://www.silastechnologies.com.

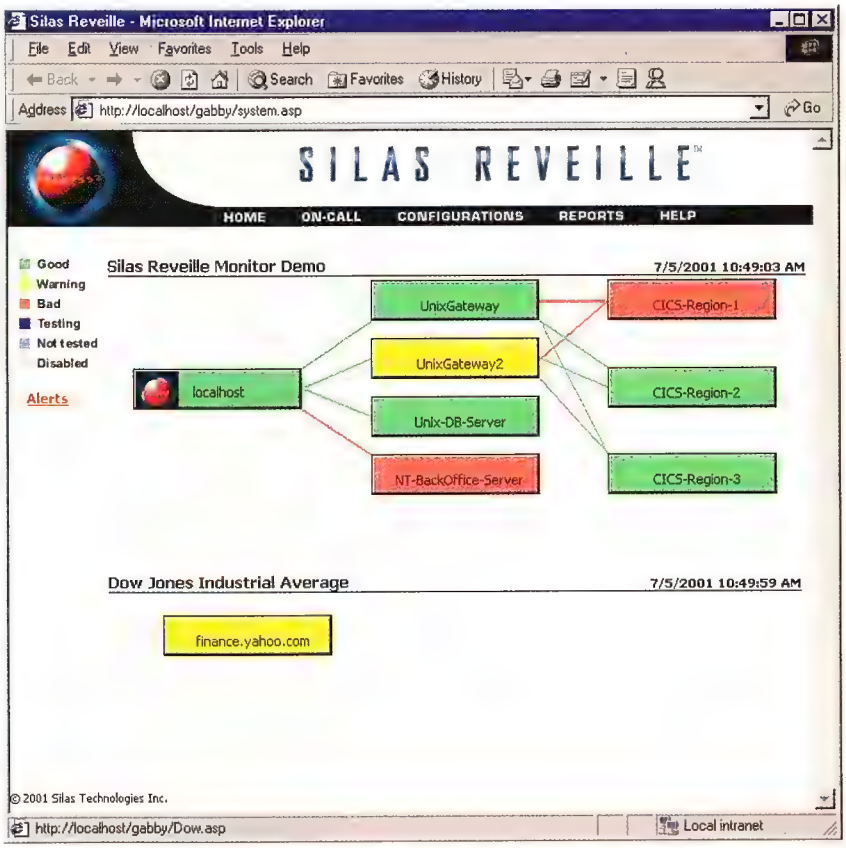

Silas Reveille 2.0 from Silas Technologies provides real-time monitoring of complex business applications and operational processes across heterogeneous computing platforms.

Please send new product announcements to products@computer.org.

ROSE-HULMAN INSTITUTE OF TECHNOLOGY

The Electrical and Computer Engineering (ECE) Department at Rose-Hulman Institute of Technology is looking for a few good people that are interested in college teaching as a vocation. The ECE Department has the tradition of emphasizing **hands-on** and **minds-on** learning through rigorous laboratory experiences designed to help students discover and understand the laws, principles, and concepts of electrical and computer engineering. Rose-Hulman has established a solid reputation for academic excellence. This reputation was built on hard work, dedication, and a clear vision "**To Be The Best**". Therefore, we invite applications for a tenure-track position in computer engineering beginning Fall 2003. Specific areas of interest are **computer architecture and software engineering**. Qualifications include an earned doctorate in electrical engineering, computer engineering, or related discipline. Again, candidates for this position must have a strong commitment to undergraduate engineering education and life-long personal professional development. Submit by October 7, 2002, a vita, a statement of teaching/research philosophy, and three references to Dr. Fred Berry, Department Head ECE, at Rose-Hulman Institute of Technology, 5500 Wabash Av., Terre Haute, IN 47803. Rose-Hulman Institute of Technology is an Equal Opportunity Affirmative Action Employer.

TOYOTA TECHNOLOGICAL INSTITUTE AT CHICAGO
Computer Science at TTI-Chicago

Toyota Technological Institute (TTI-Japan) is founding a new Department of Computer Science (TTI-Chicago) adjacent to the University of Chicago campus. Applications are invited for tenure-track and tenured faculty positions at all ranks.

TTI-Chicago will have exclusive use of the interest on a fund of $100 million being set aside by TTI-Japan for this purpose. TTI-Chicago will be dedicated to basic research, education of doctoral students, and a small masters program. Faculty members will receive continuing research grants and will have a teaching load of at most one course per year. TTI-Chicago will have close ties with the Computer Science Department of the University of Chicago.

Initial faculty appointments will commence in Autumn 2003, though some appointments may begin earlier by mutual agreement. The Department is projected to grow to a steady-state of thirty faculty by 2007.

Faculty are particularly sought with research programs in

- computational geometry
- databases and data mining
- human-computer interaction
- large-scale scientific simulation
- machine learning
- networking and distributed computing
- software and programming systems
- theoretical computer science

An advisory committee from the University of Chicago and Argonne National Laboratory will recruit the founding faculty, who will then assume leadership to determine the character of the department.

For more information, contact:
Mr. Frank Inagaki
Treasurer and Secretary to the Board
Toyota Technological Institute at Chicago
finagaki@uchicago.edu

ETH ZURICH
Professor of Computer Science (Information Technology and Education)

The new professor is responsible for research and teaching in information technology and its application to education. The design and realization of information and communication technology teaching tools is of special interest, as well as the practical testing and the didactical evaluation of new concepts. In addition, teaching concepts for life long learning and education in schools should be developed and evaluated.

The new professor is expected to be interested in teaching introductory computer science courses as well as courses in didactics of computer science and courses for non-engineering students. Applications are solicited from candidates with internationally recognized research credentials and teaching experience.

Please submit your application together with a curriculum vitae, list of publications, the names of at least three referees, and a short overview of your research interests **to the President of ETH Zurich, Prof. Dr. O. Kübler, ETH Zentrum, CH-8092 Zurich no later than October 31, 2002.** The ETHZ specifically encourages female candidates to apply with a view towards increasing the proportion of female professors.

General information about the Department of Computer Science of ETH Zurich can be found at "http://www.inf.ethz.ch/". Questions concerning this open position can be addressed to Prof. P. Widmayer, Head, Department of Computer Science, CH-8092 Zurich (e-mail: widmayer@inf.ethz.ch).

UNIVERSITY OF MARYLAND, COLLEGE PARK
Director and 7 Faculty Positions Center for Bioinformatics and Computational Biology

The University of Maryland invites faculty applications at the assistant, associate, and full professor level for the newly established Center for Bioinformatics and Computational Biology. The campus has committed resources to recruit up to eight new faculty in the Center, including a Director. It is anticipated that the primary specialization areas of the faculty will collectively span fields of computer science, mathematics and statistics, biology, and biochemistry. Their primary responsibility will be to lead a nationally visible research program in selected areas of computational genomics, proteomics, molecular evolution and phylogenetics, complementing existing strengths at the University of Maryland. Candidates for the Director position are expected to be senior researchers with prominent recognition in these areas. All other applicants are expected to have publications and research experience beyond the Ph.D.

degree with strong components of biological science and computing. Experience in interdisciplinary collaboration is an asset. The faculty will be housed in contiguous space set aside for the Center and will have access to significant high-end computing infrastructure through the University of Maryland Institute for Advanced Computer Studies. Each faculty member will also be affiliated with at least one other campus academic unit appropriate to her/his interests. There is ample potential for collaboration with other outstanding bioinformatics research groups nearby, in organizations such as NIH, Celera, TIGR, the Maryland Biotechnology Institute, and the Smithsonian Institution. To apply, send a letter of application, curriculum vitae, letters of recommendation, and URL for additional information to the search committee, in care of the appropriate departmental representative. See http://www.umiacs.umd.edu/centers/bio.htm for more information about the Center and the application procedure.

The University of Maryland is an affirmative action, equal opportunity employer. Women and minorities are encouraged to apply. Applications completed by November 30, 2002 will receive full consideration.

MISSISSIPPI STATE UNIVERSITY

Mississippi State University invites applications for several tenure-track faculty positions in the Department of Electrical and Computer Engineering; all ranks will be considered. Applicants should have a strong background in one of the following areas: (1) analog and mixed signal circuits and systems, (2) computer systems, including both hardware and software experience. A Ph.D. and a clear potential for distinguished performance in undergraduate and graduate teaching are required. Individuals will be expected to demonstrate ability to conduct sponsored research in one of the above areas and industrial experience is highly desirable. Applications will be accepted until positions are filled. Send resume and list of three references to Dr. G.M. Molen, Department of Electrical and Computer Engineering, Box 9571, Mississippi State, MS 39762. MSU is an AA/EOE.

UNIVERSITY OF CINCINNATI

The Department of Electrical & Computer Engineering and Computer Science invites applications for a Research Assistant Professor position in Computer Engineering and Computer Science. Primary responsibilities will be to conduct research in the area of wireless and mobile systems, coordinate the efforts of graduate students, and assist with teaching activities. Candidates must have an earned Ph.D. degree in Computer Science, Computer Engineering, or a closely related field. Cur-

MIT

FACULTY POSITION

Division of Health Sciences and Technology, and Department of Electrical Engineering and Computer Science

The Massachusetts Institute of Technology seeks candidates for a tenure-track faculty position offered jointly in the Harvard University-Massachusetts Institute of Technology Division of Health Sciences and Technology (HST) and the MIT Department of Electrical Engineering and Computer Science (EECS).

We seek candidates with backgrounds in interdisciplinary fields of biomedical engineering and/or the bio-medical sciences. Faculty duties will include teaching at the undergraduate and graduate levels, research, and supervision of theses. Collaborative opportunities with investigators at MIT, Harvard University, Harvard Medical School, and its teaching hospitals are abundant. In addition, with access to an exceptional student body, there are rich opportunities to build an interdisciplinary, biomedically oriented research group. There are currently 380 students enrolled in the HST's degree programs as MS, Ph.D., and/or MD candidates, and 750 students enrolled in EECS's MS and Ph.D. programs. We seek candidates who will interact productively with students and faculty in both HST and EECS, thereby fostering interdisciplinary research and teaching.

HST and EECS are strongly committed to diversity in their faculties and student populations. We encourage applications from men and women of all demographic backgrounds. All candidates should reply to the address below. The application should include curriculum vitae, statements of professional interests in both research and teaching, and the names and addresses of three or more individuals who will provide letters of recommendation. Please arrange to have such letters sent directly to the address below. The deadline for receipt of applications is November 15, 2002.

Please send applications to: **Martha L. Gray, PhD, Co-Director, Harvard-MIT Division of Health Sciences and Technology, Massachusetts Institute of Technology, E25-510, 77 Massachusetts Avenue, Cambridge, MA 02139-4307.** For more information on HST, please visit **http://hst.mit.edu.** For more information on EECS, please visit **http://www-eecs.mit.edu**

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
An Equal Opportunity/Affirmative Action Employer
Non-Smoking Environment
web.mit.edu/personnel/www

Senior Systems Architect

Specialized Supercomputer for Computational Drug Design

Successful, well-funded, rapidly growing early-stage venture based in New York City seeks extraordinarily gifted digital systems architect to assume a senior role in building a special-purpose supercomputer designed to fundamentally transform the process of drug discovery within the pharmaceutical industry.

This project aims to combine an innovative, massively parallel architecture incorporating application-specific integrated circuits with novel mathematical techniques and groundbreaking algorithmic advances in computational chemistry to direct unprecedented computational power toward the solution of key problems in molecular design. The successful candidate will be working closely with a number of the world's leading computational chemists and biologists, and will have the opportunity not only to participate in an exciting entrepreneurial venture with considerable economic promise, but to make fundamental contributions within the fields of biology, chemistry, and medicine.

The ideal candidate will be unusually intelligent and creative, with an extraordinary record of academic and/or professional achievement, and will have considerable expertise in computer organization, computer arithmetic, and numerical applications and a demonstrated ability to architect, model, design, and implement complex, high-performance hardware solutions based on ASIC and/or FPGA technologies. Experience with VLIW architectures and/or digital signal processors would be particularly useful. We are prepared to reward an exceptionally well qualified individual with above-market compensation as well as equity participation.

Please send resume, along with GPAs and standardized test scores (SAT, GRE), to *careers@schrodinger.com.*

careers@schrodinger.com **SCHRÖDINGER**

ricula vitae should be sent to Ms. Diane Phillips, P.O. Box 210030, University of Cincinnati, Cincinnati, Ohio 45221-0030, or electronically to diane.phillips@uc.edu. The University of Cincinnati is an Equal Opportunity employer and encourages applications from women and minorities.

LOUISIANA STATE UNIVERSITY
Faculty Position
(one or more tenure-track positions)
Department of Computer Science

The Department of Computer Science at Louisiana State University invites applications for one or more faculty positions, generally at the Assistant Professor level, starring either in the Spring 2003 or Fall 2003. Required Qualifications (all positions): Ph.D. in Computer Science or a closely related field; commitment to excellence in both research and teaching. Candidates in all areas are encouraged to apply; however, some areas of particular interest include high performance and scientific computing, architecture, theory, programming languages, cyber security, and operating systems. Database management and information retrieval systems are also areas of interest. Salary level will be competitive and commensurate with the qualifications of the candidate.

The University is the comprehensive university in Louisiana and participates in the Governor's Information Technology Initiative, under which a number of interdisciplinary positions will be available in scalable information infrastructure areas such as high-end computing and intelligent information and data management, biological computing, medical imaging, materials science, mobile and wireless network technologies, virtual organization and commerce, and geoinformatics.

As part of this initiative, the University has acquired a Beowulf-class Supercomputer, which will be one of the world's fastest machines, and will be dedicated to high performance computing. In addition, LSU has many other research computing facilities to deal with very large, complex problems facing scientists, engineers, and industry. The departmental research facilities can be seen on our web site at http://www.csc.lsu.edu.

Current research interests of the faculty in the department include computer networks, artificial intelligence, software engineering, robotics, computer vision and image processing, databases, information retrieval, parallel and distributed algorithms, theoretical computer science and high performance computing. The Department offers B.S., M.S., and Ph.D. degrees.

A complete application shall include curriculum vitae, list of publications, names and addresses of at least three references, and a statement of research and teaching objectives. The department will begin reviewing applications December 15, 2002 and will continue until candidate is selected. Applications and inquires should be sent to:

Professor S.S. Iyengar, Chair
Department of Computer Science
Louisiana State University
Ref: Log #0001
Baton Rouge LA 70803-4020
(225) 578-1252
iyengar@bit.csc.lsu.edu
LSU IS AN EQUAL OPPORTUNITY/
EQUAL ACCESS EMPLOYER

THE ALBERT-LUDWIGS-UNIVERSITY FREIBURG

The Albert-Ludwigs-University Freiburg is establishing a new interdisciplinary Center for Bioinformatics (ZfBI) with three Professorships. This Center is embedded in the university-wide research focus in the area of life sciences.

The Faculty of Applied Sciences, with its institutes of Computer Science and Microsystem Technology, invites applications for the position of a

Full Professor (C4) in Bioinformatics

within the new Center. The prospective candidate will belong to the Institute of Computer Science and will also be an

---

PENNSTATE
University Park

IT RESEARCH FACULTY AND ENGINEERS

The Applied Research Laboratory at Penn State performs fundamental and applied research in a broad range of disciplines. ARL's Information Science and Technology (IST) Division is conducting major research projects in the fields of distributed sensing, networked robotics, system failure diagnostics, emergent collaboration, mobile code technologies, discrete event control, autonomous systems, and information brokering. To support its multi-university research teams, the IST Division has position openings in the areas of INTELLIGENT CONTROL RESEARCH, DISTRIBUTED SYSTEMS SOFTWARE, NETWORKED SYSTEMS, and SCIENTIFIC SOFTWARE DEVELOPMENT. To apply and for a description of position and information about the Applied Research Laboratory you are invited to visit our website at www.arl.psu.edu/index.html and click on Employment. Please indicate (Source of Ad) in your application cover letter. Please contact us at arl-jobs@psu.edu if you have any questions or need assistance. U. S. CITIZENSHIP REQUIRED.

Penn State is committed to affirmative action, equal opportunity and the diversity of its workforce.

Faculty Position in Bioinformatics and Computational Biology
Departments of Biological Sciences and Biomedical Engineering
in conjunction with the
Roy J. Carver Center for Comparative Genomics
and The Coordinated Laboratory for
Computational Genomics at
THE UNIVERSITY OF IOWA

A tenure-track position at the Assistant Professor level is open for an individual interested in computational methods and their application to significant biological problems. Areas of particular interest include: comparative genomics, genome evolution, genomic applications, computational methods, and algorithm design; however, applicants in other related areas will also be carefully considered. This position is part of a campus-wide initiative involving many departments in the Colleges of Liberal Arts and Sciences, Engineering, Medicine, and Public Heath addressing all aspects of Bioinformatics. The successful candidate will encounter a diverse and dynamic community of scholars interested in computational methods and their application to biological problems.

Candidates will be expected to have a Ph.D. as well as a record of accomplishment as reflected in publications in leading journals. The successful candidate will be expected to establish and maintain an extramurally-funded research program. Recently renovated space and a competitive start-up package will be made available. Applicants should send a curriculum vitae, statement of research objectives, selected reprints, a description of teaching interests, and the names of three references to:

Computational Biology Search Committee
The University of Iowa
Department of Biological Sciences
138 Biology Building
Iowa City 52242-1324

Applications will be considered as they are received. *The University of Iowa is an affirmative action/equal opportunity employer. Women and minorities are especially encouraged to apply.*

associated member in the Faculty of Biology. He/She will lead research and teaching in the field of Bioinformatics. He/She will also be in charge of implementing new courses and degrees in cooperation with other universities in the upper Rhine valley area network (EUCOR).

Applications, including a curriculum vitae, publications list and a letter stating the research interests should be sent by October 4, 2002 to the Dean of the Faculty of Applied Sciences, University of Freiburg, D-79110 Freiburg, Germany. http://www.faw.uni-freiburg.de/.

At the Faculty of Biology, within the framework of the establishment of the ZfBI, a Professorship (C3) for Experimental Bioinformatics and a third Professorship (C4) in Bioinformatics and Molecular Genetics are in the process of appointment.

The University of Freiburg aims to increase the percentage of women in research and teaching, and therefore encourages female candidates meeting the above qualifications to apply.

UNIVERSITY OF NEW ORLEANS

The Department of Computer Science invites applications for two tenure track positions. The appointments will be primarily at the rank of assistant professor with very competitive salary/start up funds and excellent benefits in a very affordable cost of living area.

One of the two positions is in Bioinformatics and is jointly funded by the University of New Orleans and the Research Institute for Children. We are particularly interested in applicants with a research interest that can be applied to problems in microbial informatics, such as comparative genomics, gene expression, or predictions of functional activities. Qualifications for this position include a PhD. in Computer Science, Bioinformatics, or a closely related field. Post-doctoral experience and substantial knowledge and expertise with biological systems are desirable.

For the second position, the department has a particular interest in specialists in systems, networks, computer security, databases, and software engineering, but will give serious consideration to excellent applicants from any of the computing fields. Qualifications for this position include a Ph.D. in Computer Science or a closely related field, and strong practical skills.

Responsibilities for both positions include teaching at the B.S., M.S., and Ph.D. levels, supervision of graduate students, securing external funding, and publication of research results. Our computing facilities include state-of-the-art instructional and research laboratories. This includes a 72-node Beowulf High Performance class system dedicated to research. Excellent opportunities for interaction with federal agencies and corporations located at the University of New Orleans' (UNO) Research & Technology (R&T) Park, and the Stennis Space Center in nearby Mississippi exist. Our newly approved Canizaro/Livingston Information Management & Bioinformatics (CLIMB) Center (located at UNO's R&T park) provides additional opportunities for research funding.

The University of New Orleans enrolls over 17,000 students, is an urban public university, and is a member of the Louisiana State University system. The Computer Science Department has over 700 undergraduate majors, and its baccalaureate degree program is accredited by the Computing Commission of ABET. The Department initiated a masters degree program in Fall 1991, which has since grown to some 80 students. The Department also participates in a multidisciplinary doctoral program in engineering and applied science.

The Research Institute for Children is a newly established, well-funded, academic collaboration based at Children's Hospital of New Orleans. A 60,000 sq. ft state-of-the-art research building was dedicated in February of this year. The research focus is host-pathogen interactions in the causation of human disease.

Applicants should send a resume and the names of at least three references to: Search Committee, Department of Computer Science, University of New Orleans, New Orleans, LA 70148. E-mail: search@cs.uno.edu. Applications will be accepted until the positions are filled. Women and minorities are especially encouraged to apply. The University of New Orleans is an Equal Opportunity/Affirmative Action Employer.

CLEMSON UNIVERSITY

Applications are invited for faculty positions in the Computer Engineering area of the Department of Electrical and Computer Engineering. The Department has strong research programs in wireless communications and signal processing, mechatronics, computational electromagnetics, solid-state device reliability, power systems, cluster-based computing, reconfigurable computing, and machine vision.

Exceptional candidates at all levels and in all research areas related to computer engineering will be considered. However, we intend to fill at least one position in the area of computer communication networks. For the other positions, we are interested in individuals who can contribute to the Department's active research programs or who can serve as conduits for building interdisciplinary research teams in emerging areas at Clemson (e.g., robotics/control, intelligent systems, and computer simulation).

Candidates should hold a Ph.D. degree in Computer Engineering, Electrical Engi-

---

BALL STATE UNIVERSITY

Ball State University

Muncie, Indiana

Assistant Professor

Department of Computer Science

The Department of Computer Science seeks applicants for a full-time, tenure-track faculty position at the Assistant Professor rank, available August 22, 2003. Responsibilities include but are not limited to: teaching a wide variety of undergraduate and graduate courses in computer science; an active productive research program; serving on departmental committees.

Minimum qualifications: master's degree in computer science with all coursework completed toward a doctorate in computer science; 15 or more years of teaching graduate and undergraduate computer science classes; experience teaching a wide range of upper division and graduate classes, including software engineering, programming languages, operating systems, and compiler construction; a proven track record of grantsmanship, including federal, state, and industrial support; expertise in software design; extensive experience in supervising graduate student thesis work; published in respected journals in the field; shown that his/her research has visibility beyond academe and has been put into practice through technology transfer activities; presented research results in international and national conference settings; experience and demonstrated skill at interacting with industry; demonstrated familiarity with patent and copyright procedures at a university.

Ball State University has approximately 18,000 students. The Department of Computer Science has approximately 200 undergraduate majors and 100 M.S. students. Departmental lab facilities include both Microsoft Windows based machines and Unix based machines. For more information, visit web pages at: www.bsu.edu/ and www.cs.bsu.edu/.

Send letter of application; vita; copies of transcripts; and the names, addresses, and telephone numbers of at least three references to: Dr. Paul Buis, Computer Science Search Committee, Department of Computer Science, Ball State University, Muncie, IN 47306. Transmission of materials by e-mail is encouraged at cssearch@cs.bsu.edu. Review of completed applications will begin immediately and will continue until the position is filled. (www.bsu.edu)

Ball State University is an equal opportunity, affirmative action employer and is strongly and actively committed to diversity within its community.

neering, Computer Science, or a closely related field and should have high potential for establishing a sustained research program and quality teaching. The individual selected will be expected to contribute to both new and ongoing research programs at Clemson and to teach both undergraduate and graduate courses. A detailed description of the department is available at http://www.ece.clemson.edu. Send resume and names and addresses of five references to Chair, Holcombe Department of Electrical and Computer Engineering, 105 Riggs Hall, Box 340915, Clemson University, Clemson, SC 29634-0915. Evaluation will begin October 1, 2002, and will continue until the positions are filled. Clemson University is an Equal Opportunity/Affirmative Action Employer.

INFORMATION SYSTEM MANAGER

Design & Implement Local & Wide Area Networks. Implement and manage corporate network, email and websites including backup management and firewall configuration. Bachelors Degree in Business Administration or MIS required. Send ad w/resume to: Universal Dealer Consultants; 29B Technology Drive, Suite 150; Irvine, CA 92618.

ETH ZURICH
Professor in Software Technology

The successful applicant for this open position is expected to be an internationally recognized expert in the software field, with a strong publication record and extensive practical experience. Although excellent candidates will be considered regardless of their area of specialization in software technology, fields of particular interest include: component technology; object technology; formal methods and their practical applications; programming languages for software engineering; software validation and verification; management of software projects. The Professor will be invited to join initiatives in progress such as a Component Certification Center.

All professors are expected to participate in teaching introductory computer science courses.

ETH has one of the foremost computer science departments in Europe with a prestigious tradition of research and system development in the fields of programming language, software engineering and component development.

Please submit your application together with a curriculum vitae, list of publications, the names of at least three referees, and a short overview of your research interests **to the President of ETH Zurich, Prof. Dr. O. Kübler, ETH Zentrum, CH-8092 Zurich no later than October 31, 2002**. The ETHZ specifically encourages female candidates to apply with a view towards increasing the proportion of female professors.

General information about the Department of Computer Science of ETH Zurich can be found at "http://www.inf.ethz.ch/". Questions concerning this open position can be addressed to Prof. P. Widmayer, Head, Department of Computer Science, CH-8092 Zurich (e-mail: widmayer@inf.ethz.ch) or Prof. Dr. B. Meyer (bertrand.meyer@inf.ethz.ch).

GEORGIA STATE UNIVERSITY
Department of Computer Science

The Department of Computer Science of Georgia State University invites applications for an anticipated tenure-track position for assistant/associate professor beginning the Fall semester, 2003. Earned Ph.D. in Computer Science, or a closely related discipline, and an excellent record in publications in computer science are required with preference for extramural funding. Preference is for individual with specialty in software engineering, graphics, content processing, internet multimedia communications, operating systems, or algorithms.

The department offers programs leading to the B.S., M.S., and Ph.D. degrees in computer science. Departmental computing facilities for research and instruction include a departmental network of PCs, Unix/Linux workstations, and a 24-processor Origin 2000 high-performance computer and five laboratories, and a hypermedia and visualization research laboratory. A full-time systems programmer supports the departmental computing facilities.

Applicants should send a letter of application, vita without birthdate, but with citizenship status, and three letters of reference and transcripts of all graduate and undergraduate work to: Chair, Department of Computer Science, Georgia State University, University Plaza, Atlanta, Ga. 30303-3083 (or e-mail to: mfraser@cs.gsu.edu). Applications will be accepted until position is filled. Georgia State University is an EEO/AA institution.

SAN FRANCISCO STATE UNIVERSITY

The Department of Computer Science invites applications for the position of department Chair to begin by August 22, 2003. We are seeking candidates with outstanding records in both leadership and research. Administrative experience is also desired. A successful candidate should have a Ph.D. in Computer Science in a software-related field and be qualified for an appointment at the level of Full Professor. The new Chair will provide leadership in the expansion of departmental research activities, revision of the Master of Science program, professional development of junior faculty, and expansion of industrial connections and funded research. In addition, the new Chair will teach undergraduate and/or graduate courses with a reduced load. The Department of Computer Science (http://cs.sfsu.edu) offers an ABET-accredited undergraduate program and a Master of Science degree. We serve approximately 1150 students. Over the past 3 years, the student body has grown nearly 10% each year. We have 11 full-time faculty as well as several who serve part-time. SFSU is an Affirmative Action Employer and women and minorities are encouraged to apply. Interested candidates should send a curriculum vitae, statements of research, teaching, and leadership plans, the names of at least three references, and copies of representative publications to Chair Search Committee, Computer Science Department, San Francisco State University, 1600 Holloway Avenue, San Francisco, CA 94132. Review of applications will begin November 4, 2002 and will continue until the position is filled.

CLEVELAND STATE UNIVERSITY

The Department of Electrical and Computer Engineering at Cleveland State University invites applications for a tenured/tenure-track faculty position that will lead the development of the Master of Science in Software Engineering degree program. The appointment will be at the rank of **Associate/Full Professor**, with competitive salary commensurate with experience. Applicants must document relevant prior teaching and research experience in software engineering and the potential for continuing research support. Major local corporations and government agencies, such as Rockwell Automation, Philips, General Electric, Computer Associates, ABB Automation, Keithley Instruments, Nordson, Parker Hannifin, Eaton Corp., NASA Glenn Research Center, and many others, have a very strong need for software engineers and have expressed their desire to support the development of the MSSE program and share their expertise. Minimum qualifications: Ph.D. in Software Engineering, Computer Engineering, Computer Science or related fields with specialization in software engineering; teaching experience; record of publications; and record of funded research. Preferred qualifications: distinguished record of teaching experience in software engineering; distinguished record of publications in a recognized area of software engineering; sustained record of externally funded research; experience in laboratory development; excellent communication skills; relevant industrial experience. A letter of application and current resume, including names, addresses, phone numbers and e-mail addresses of at least three references, should be sent to: **Search Committee, Electrical and Computer Engineering Department, Cleveland State University, Cleveland, Ohio 44115-2425, phone (216) 687-2589. For information about CSU visit our home page: www.csuohio.edu**

Applications will be accepted until the position is filled. Cleveland State University is an equal opportunity, affirmative action employer. Women, persons with disabilities and members of minority groups are encouraged to apply.

BYU

Applications are invited for tenure-track faculty positions. Applicants must have a Ph.D. and strong research orientation to support the Ph.D. and M.S. programs. For more information, see http:www.cs.byu.edu. Send curriculum vitae to E. Daniel Johnson, 3362 TMCB, BYU, Provo, Utah 84602. BYU is an EEO/AA employer with preference for LDS candidates.

WRIGHT STATE UNIVERSITY
Assistant, Associate, Full & Non-Tenure Track Faculty Positions

The Department of Computer Science and Engineering at Wright State University seeks applicants for tenure track Assistant, Associate or Full Professor positions and for non-tenure track positions as Instructor or Lecturer. The Department resides in the College of Engineering and Computer Science and offers B.S., M.S. and Ph.D. degrees both in Computer Science and Computer Engineering. Candidates for these positions are expected to have an earned Ph.D. (except for the Instructor/Lecturer position for which an M.S. and admissibility to our Ph.D. program are required) in computer science, computer engineering, or a closely related field and evidence of scholarship in that field appropriate to the position. Associate and Full Professor applicants should have a distinguished record in computing, demonstrating strong leadership in both research and teaching, commensurate with the rank. There is specific interest in faculty specializing in networking, operating systems, particularly in parallel and real time applications, and in bioinformatics. However, all high quality applicants will be considered regardless of their field of specialization.

The Computer Science and Engineering Department is one of four departments in the College of Engineering and Computer Science. The Department currently has 24 faculty members, 525 undergraduate majors, and more than 125 M.S. and 35 Ph.D. students. It is housed in a new, attractive, engineering building with fully networked Unix and NT environments, an Origin 2000 supercomputer, an NCR4800 Teradata machine, and excellent research laboratories. Wright State University, an institution of 16,000 students, is located on a spacious campus with a significant area of protected green space in a growing high-technology suburban community. It is surrounded by commercial (NCR, Lexis-Nexis, Reynolds & Reynolds, Mead, etc.) and government (WPAFB) research and development facilities. The University is proactively committed to industrial and government partnerships for research and economic development ventures and has established the Information Technology Research Institute closely aligned with the CS&E Department. The Department receives special support for enhancement of its graduate program from WSU, and the Ohio Board of Regents. Last year, the Department received $5.3M in externally sponsored research. A variety of affordable and pleasant living environments, schools and parks, attractive to professionals and their families, are conveniently located close to the campus.

Applicants should clearly indicate the position for which they are applying, including the rank sought, and provide a brief statement of their research and teaching interests and goals. They should include a complete vitae with names, addresses, telephone numbers and e-mail addresses of at least three references, documentation of teaching abilities, plus any additional supporting information that they would like to provide. Salaries are highly competitive. Address applications and supporting information to:

Chair, Faculty Search Committee
Department of Computer Science & Engineering
Wright State University
Dayton, OH 45435

Consideration of candidates starts September 16, 2002 and continues each month until September 30, 2003, or until the positions are filled. For details and information you may call (937) 775-5134 or contact Oscar N. Garcia, NCR Distinguished Professor and Chair, ogarcia@cs.wright.edu

Wright State University is an equal opportunity/affirmative action employer.

STANFORD UNIVERSITY
Department of Computer Science
Faculty Openings

The Computer Science Department of Stanford University invites applications for tenure-track faculty positions at both the junior level (Assistant or untenured Associate Professor) and senior (tenured Associate or Full Professor) level. We are seeking applicants from all areas of Computer Science, including Foundations, Artificial Intelligence, Graphics, Databases, Systems, Human Computer Interaction, and Networking. The department also has interest in applicants doing research at the frontiers of computer science, for instance biological computing, bio-informatics, computation and arts, or computational economics. Higher priority will be given to the overall innovation and promise of the candidate's work than to any specific

Computer Science Department Chair

Union seeks applications and nominations for a faculty appointment as Chair of the Computer Science Department for September 2003. This position will be added to the five full-time positions now in the department. The department offers a BS degree and has a small MS program. It shares additional faculty members with the Electrical and Computer Engineering Department, with which it offers a Computer Engineering major. We are seeking a dynamic person who can lead a growing department in maintaining and updating a strong major while embracing a new College-wide "Converging Technologies" initiative focused on teaching new technologies through interdisciplinary cooperation in engineering, sciences, arts and humanities. Candidates should hold a Ph.D. in computer science or a closely related discipline, and be committed to excellence in teaching and research.

Union is a highly selective coeducational college with programs in the liberal arts and engineering, and with a strong tradition of undergraduate research. The department offers a congenial working environment, excellent students, and up-to-date laboratories with Unix workstations and PCs.

Union College is located in New York's Capital Region, three hours from New York City and Boston. Rensselaer Polytechnic Institute, the University at Albany, Knolls Atomic Power Laboratory, and the GE Global Research Center are nearby and offer unique opportunities for collaborative teaching and research.

Salaries and fringe benefits are competitive. For more information see **www.cs.union.edu**. Applications should include names, postal and email addresses of three references, and a statement of teaching and research philosophy for an undergraduate computer science department to:

Search Committee Chair
Computer Science Department
Union College, 807 Union Street, Schenectady, NY 12308

Review of applications will start in November and continue until the position is filled.

Union College is an equal opportunity, affirmative action employer and is strongly committed to increasing the diversity of its employees.

area.

An earned Ph.D., evidence of the ability to pursue a research program, and a strong commitment to graduate and undergraduate teaching are required. For senior appointments, we are searching for strong, energetic and visionary leaders. Successful candidates will be expected to teach computer science courses at the graduate and undergraduate levels and to build and lead a team of graduate students in Ph.D. research. Further information about the Computer Science Department can be found at http://cs.stanford.edu.

Applications should include a curriculum vita, statements of research and teaching interests and the names of at least four references. Candidates are requested to ask references to send their letters directly to our search committee. The letters should be sent in as soon as possible, but no later than the application deadline. All materials should be sent to: Search Committee Chair, c/o Laura Kenny-Carlson, Computer Science Department, Stanford University, Gates 278, Stanford, CA 94305-9025; or via electronic mail to search@cs.stanford.edu.

The review of applications will begin on January 6, 2003, and applicants are strongly encouraged to submit applications by that date; however, applications will continue to be accepted until February 3, 2003. The positions are available beginning Autumn 2003. Stanford University is an equal opportunity employer and welcomes nominations of women and minority group members and applications from them.

UNIVERSITÄT KARLSRUHE (TH)
Computing Center
(http://www.rz.uni-karlsruhe.de)

At the University of Karlsruhe, the oldest technical university in Germany, the distributed computing environment for information processing and digital communication has been continuously extended under the control of the Computing Center. For the management of this large area applications for the following professorship are invited:

Associate Professorship (C3) in the field of Distributed Systems and Computer Network Services

We are looking for a personality with relevant expertise as well as the ability to represent this field appropriately in research and teaching. Thus the following aspects are of special interest:

- Networking and Communication Technologies
- Management of Distributed Systems
- Security on Networks and Distributed Systems
- Multimedia Network Services and appropriate Infrastructures
- Integration of Information Systems

As the successful candidate will also assume executive functions at the Computing Center, substantial experience in the management of personnel as well as scientific projects is expected beside an eminent scientific qualification.

Habilitation (standard qualification for a professorship in Germany) or an equivalent qualification is required. A basic command of German is expected.

In case of a first-time call to a professorship the employment status is basically limited in time, exceptions to the limitation are possible.

The University of Karlsruhe encourages especially qualified women to apply. Handicapped persons with the same qualifications will be given preference.

Applicants are asked to submit a letter of application, together with a curriculum vitae, a list of publications and a summary of their professional career as well as details on their teaching experience to:

Prof. Dr. Wilfried Juling
Rechenzentrum der Universität
Karlsruhe (TH)
D-76128 Karlsruhe
Telephone: ++49/(0)721/608-3158
Fax: ++49/(0)721/32550
E-mail: juling@rz.uni-karlsruhe.de

The closing date is 4 weeks after publication.

PROGRAMMER ANALYST

Develop multi tier applications using various development tools & perform all steps required during the life cycle of a project. BS in Computer Science, Engineering or related field required. Send ad w/ resume to: OSI Consulting, Attn: MC, 5950 Canoga Ave. #300, Woodland Hills, CA 91367.

INFORMATION AND COMMUNICATIONS UNIVERSITY
Full-time and Visiting Faculty Positions

Information and Communications University (ICU) invites applicants, regardless of nationality, for full-time and visiting faculty positions in all areas of computer science and engineering including computer graphics, HCI, robotics, natural language, bioinformatics, software engineering, computer networks, and information security. Ph.D. in computer science/engineering or related fields and evidence in teaching and research excellence are required. ICU is a fast growing university aiming to educate leaders in the information industry and conduct top quality research. ICU offers BS, MS and Ph.D. degree programs in computer science/engineering, telecommunications, electronics engineering, and management of technology. ICU has opened an undergraduate program this spring and successfully recruited students who are ranked at top 2% of 2001 high school graduates. The university is located in Taeduk science town, Taejeon, Korea. Please refer to the university home page (http://www.icu.ac.kr) for further information. Salaries are highly competitive and commensurate with qualifications and experience. Successful candidate should have his/her willingness to give lectures in English since the teaching language is English. Positions start Spring 2003. The positions will remain open until filled. Teaching or industrial experience will be highly desirable. To apply, please send e-mail a resume, a statement of teaching and research interests, and three references to:

Prof. Myungchul Kim, Dean,
School of Engineering
Information and Communications University
139 Kajung-Dong, Yusung-Ku, Taejon, KOREA
E-mail: mckim@icu.ac.kr
Tel: +82-42-866-6100,
Fax: +82-42-866-6110

ETH ZURICH
Professor of Computer Science (Theoretical Computer Science)

The new professor is responsible for research and teaching in the broad area of theoretical computer science. Any research specialization within theory and its application is welcome.

The professor shall lead a research group and cooperate with other research groups in the department and/or in related disciplines at ETH (e.g. mathematics or biology). Candidates are expected to show internationally recognized research credentials and proven teaching ability.

Please submit your application together with a curriculum vitae, list of publications, the names of at least three referees, and a short overview of your research interests **to the President of ETH Zurich, Prof. Dr. O. Kübler, ETH Zentrum, CH-8092 Zurich no later than October 31, 2002.** The ETHZ specifically encourages female candidates to apply with a view towards increasing the proportion of female professors.

General information about the Department of Computer Science of ETH Zurich can be found at "**http://www.inf.ethz.ch/**". Questions concerning this open position can be addressed to Prof. P. Widmayer, Head, Department of Computer Science, CH-8092 Zurich (e-mail: widmayer@inf.ethz.ch) or Prof. U. Maurer (e-mail: maurer@inf.ethz.ch).

ST. MARY'S UNIVERSITY
Engineering Department
Two Faculty Positions

The Department of Engineering at St. Mary's University of San Antonio invites applications for two tenure track faculty positions, one in Computer Engineering and another in Electrical Engineering in the area of Control Systems at the Assistant Professor levels. Applicants must have earned a doctorate degree and must demonstrate strong potential for excellence in teaching and research. The suc-

cessful candidates are expected to teach a variety of courses at both undergraduate and graduate levels. For a complete description of the positions, please see: http://www.engr.stmarytx.edu/faculty_positions. Positions are open until filled. Review of applications will begin immediately. Applicants should forward a resume including a statement of goals and a plan for teaching and research, and names, addresses, and telephone numbers of at least three references to: Dr. Djaffer Ibaroudene, Search Committee Chair, St. Mary's University, One Camino Santa Maria, Box 34, San Antonio, TX 78228. Phone: 210-431-2050, e-mail: dibaroudene@stmarytx.edu. St. Mary's University is an Affirmative Action and equal opportunity employer.

UNIVERSITY OF MINNESOTA - TWIN CITIES
Faculty Positions
Department of Computer Science and Engineering

The Department of Computer Science and Engineering at the University of Minnesota invites applications for several open faculty positions. These positions are primarily at the assistant professor level, but highly qualified applicants at higher ranks will also be considered. Specialists from all areas of computer science and engineering are encouraged to apply. Requirements include a Ph.D. in computer science, computer engineering, or a closely related discipline, a commitment to quality teaching, and the potential for carrying out outstanding research. Senior candidates must possess a distinguished record of teaching, research, and service.

The research and teaching missions of the Department encompass a wide range of areas, including networking and distributed multimedia, internet technologies, cryptography and security, databases, data mining, software engineering, computer architecture, compilers, programming languages, artificial intelligence, robotics and computer vision, computer graphics and visualization, algorithms and complexity theory, geometric computing and applications, human-computer interaction, parallel computing, and scientific computing. Its faculty members, numbering thirty-four, have access to outstanding computing facilities both within the Department and at the various research center on campus, including the Army High Performance Computing Research Center, the Minnesota Supercomputing Institute, and the Institute for Mathematics and its Applications. The Department is also a key player in the recently established Digital Technology Center at the University. External research funding in the department has grown steadily over the past several years and recently topped $4.4 Million. Additional information about the Department is available at its World Wide Web home page: <http://www.cs.umn.edu>. The Minneapolis-St. Paul area is a major center for the computer industry and for advanced technology, and the Department enjoys strong interactions with local industries.

Applicants should submit a curriculum vitae, a research summary, and the names of at least three references to:

Chair, Faculty Recruiting Committee
Department of Computer Science and Engineering
University of Minnesota
4-192 EE/CS Building
200 Union Street S.E.
Minneapolis, MN 55455

Electronic submissions of applications are welcome and may be sent via e-mail to applications@cs.umn.edu. (Electronic submissions must be in postscript, PDF or Word formats.) Review of completed applications will begin December 1, 2002, but the search will remain open until all positions are filled.

The University of Minnesota is an equal opportunity educator and employer.

UNIVERSITY OF ALABAMA AT BIRMINGHAM (UAB)
Department Chair Position
Department of Computer and Information Sciences
http://www.cis.uab.edu/

Applications and nominations are being accepted for the position of Department Chair. Candidates should possess proven leadership and administrative skills and have an established reputation as a scholar. The Chair will be expected to provide vigorous leadership in the development of interdisciplinary collaboration initiatives in bioinformatics with the highly ranked UAB biomedical research community. The successful candidate will be appointed to a tenured position as Professor within the Department, with salary at a level commensurate with qualifications and experience. Letters of application should include a full curriculum vitae and names and addresses of five references. Applications will be accepted until the position is filled; screening will begin immediately. Applicants from minority and women applicants are especially encouraged.

Nominations or applications should be sent to:

Chair_search@cis.uab.edu
or
Department Chair Search Committee
Department of Computer and Information Sciences
University of Alabama at Birmingham
Birmingham, AL 35294-1170

The University of Alabama at Birmingham is an equal opportunity/affirmative action employer.

SOFTWARE ENGINEER

Develop, debug software for storage network processor verification. Interface with ASIC, simulation engineers to identify and resolve hardware design issues. MSEE, relevant TCP and networking protocols experience with embedded processor. iSCSI and storage protocols required. Verilog simulation, debug experience, C language, GNU tools, Linux, Solaris. Astute Networks, 16868 Via Del Campo Ct. Suite 200, San Diego, CA 92127, jobs@astutenetworks.com, Reference #SE006

September 2002

NEXT ISSUE

Web Intelligence

Web Intelligence explores the fundamental role and practical impact of artificial intelligence and advanced information technology on the next generation of Web-empowered products, systems, services, and activities. We look at topics such as Web information management, information retrieval, agents, Web mining and farming, and emerging Web-based applications.

Advertising Personnel

Marion Delaney
IEEE Media, Advertising Director
Phone: +1 212 419 7766
Fax: +1 212 419 7589
Email: md.ieeemedia@ieee.org

Marian Anderson
Advertising Coordinator
Phone: +1 714 821 8380
Fax: +1 714 821 4010
Email: manderson@computer.org

Sandy Brown
IEEE Computer Society,
Business Development Manager
Phone: +1 714 821 8380
Fax: +1 714 821 4010
Email: sb.ieeemedia@ieee.org

Debbie Sims
Assistant Advertising Coordinator
Phone: +1 714 821 8380
Fax: +1 714 821 4010
Email: dsims@computer.org

Advertising Sales Representatives

Mid Atlantic (product/recruitment)
Dawn Becker
Phone: +1 732 772 0160
Fax: +1 732 772 0161
Email: db.ieeemedia@ieee.org

Midwest (product)
David Kovacs
Phone: +1 847 705 6867
Fax: +1 847 705 6878
Email: dk.ieeemedia@ieee.org

Northwest (product)
John Gibbs
Phone: +1 415 929 7619
Fax: +1 415 577 5198
Email: jg.ieeemedia@ieee.org

Southern CA (product)
Marshall Rubin
Phone: +1 818 888 2407
Fax: +1 818 888 4907
Email: mrubin@westworld.org

Southwest (product)
Royce House
Phone: +1 713 668 1007
Fax: +1 713 668 1176
Email: rh.ieeemedia@ieee.org

Southeast (product/recruitment)
C. William Bentz III
Email: bb.ieeemedia@ieee.org
Gregory Maddock
Email: gm.ieeemedia@ieee.org
Sarah K. Huey
Email: sh.ieeemedia@ieee.org
Phone: +1 404 256 3800
Fax: +1 404 255 7942

Europe (product)
Hilary Turnbull
Phone: +44 131 660 6605
Fax: +44 131 660 6989
Email: impress@impressmedia.com

New England (product)
Jody Estabrook
Phone: +1 978 244 0192
Fax: +1 978 244 0103
Email: je.ieeemedia@ieee.org

New England (recruitment)
Barbara Lynch
Phone: +1 401 738 6237
Fax: +1 401 739 7970
Email: bl.ieeemedia@ieee.org

Connecticut (product)
Stan Greenfield
Phone: +1 203 938 2418
Fax: +1 203 938 3211
Email: greenco@optonline.net

Midwest (recruitment)
Tom Wilcoxen
Phone: +1 847 498 4520
Fax: +1 847 498 5911
Email: tw.ieeemedia@ieee.org

Northwest (recruitment)
Mary Tonon
Phone: +1 415 431 5333
Fax: +1 415 431 5335
Email: mt.ieeemedia@ieee.org

Southern CA (recruitment)
Tim Matteson
Phone: +1 310 836 4064
Fax: + 1 310 836 4067
Email: tm.ieeemedia@ieee.org

Japan
German Tajiri
Phone: +81 42 501 9551
Fax: +81 42 501 9552
Email: gt.ieeemedia@ieee.org

Computer
10662 Los Vaqueros Circle
Los Alamitos, California 90720-1314
Phone: +1 714 821 8380 • Fax: +1 714 821 4010
http://computer.org • advertising@computer.org

An Economically Scalable Internet

Abdelsalam "Solom" Heddaya
InfoLibria

Over the past decade, many observers have claimed that the Internet brings the information revolution's components together in a way that will rival the industrial revolution's effects on human productivity and quality of life. Replication is central to this comparison: The industrial revolution enabled large-scale replication of increasingly complex physical objects, and the information revolution—primed by the Internet—fulfills a similar function for information objects. Yet this comparison highlights a key Internet deficiency: Economically, it does not scale well.

In manufacturing technology, the per-unit cost decreases as the total production rate increases. Successful products create a virtuous cycle—the more popular they get, the cheaper and more profitable they become. This cycle emerges from an industrial architecture that becomes more efficient in per-unit resource costs as the resources required for production increase. Doubling production resources more than doubles production capacity.

Transportation and other delivery systems evolved to match this economy of scale as these production efficiencies made their way to the marketplace. Success in this environment rewarded both the producer with higher profits and the consumer with lower prices.

BANDWIDTH RESOURCES

But the Internet's architecture does not scale economically this way. The recurring costs of Internet bandwidth—not the capital costs of hardware and software—tend to dominate the costs of Web hosting. Roughly speaking, a Web site that requires 500 Mbps costs 10 times more than one requiring 50 Mbps. In this business environment, a site's popularity does not translate into producer profits and consumer savings.

The Internet has proven its system scalability but has yet to demonstrate significant economic scalability.

Thus, from 1980 to 2000, Internet backbone traffic grew by a factor of 100 million, but the bandwidth cost for serving a single streaming advertisement in 2000 remained higher than the revenue it generated.

Similarly, the bandwidth cost of delivering one hour of a 1-Mbps video stream was roughly US$5—far more than service providers could expect consumers to pay, even though much cheaper than long-distance telephone network bandwidth. More significantly, the delivery costs applied almost equally to popular and unpopular content.

When investors saw the Internet's popularity soar among the general population in the 1990s, they wanted to invest in resources that they expected would make the Internet incredibly cheap—at least for popular content and services. Examples of such investments include the deployment of huge data centers and massive long-distance fiber capacity. There is little evidence, however, that these investments realized fundamental economies of scale, even for successful Internet services.

QoE VERSUS QoS

Nor did the end users' *quality of experience* reach a high enough or consistent enough level to convince them to pay for new services. Internet QoE remains low and unreliable. I use the term QoE rather than quality of service because QoS is not necessary for QoE, and QoE is sufficient for successful service. For example, downloading a two-hour movie to a home entertainment system or PC for later viewing can yield high QoE for the movie-viewing experience, even if the available bandwidth is of low quality in terms of packet loss, packet delay, and bandwidth variation.

To complicate matters, the Internet's commercialization and rapid growth during the 1990s spawned many scaling, quality, and security solutions that ran afoul of some dearly held architectural principles of the Internet community. The Internet Engineering Task Force—the key standards-setting body entrusted with maintaining the network's integrity—became the forum for many heated debates.

The architectural debate is best captured by Marjory S. Blumenthal and David D. Clark in "Rethinking the Design of the Internet: The End-to-End Arguments vs. the Brave New World" (*ACM Trans. Internet Technology*, vol. 1, no. 1, Aug. 2001, pp. 70-109). Regardless of the merits of

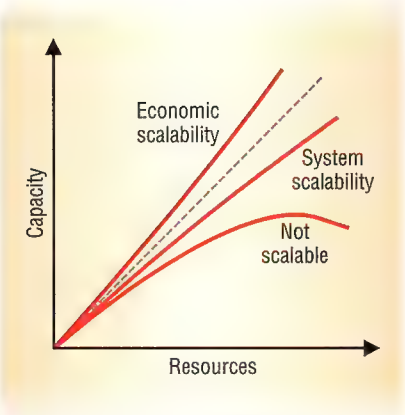

Figure 1. Economic scalability must improve the performance-price ratio superlinearly.

particular proposals or good intentions of all parties, the intensity of these debates may well have detracted from the community's ability to focus on fundamental issues.

SCALABILITIES

We can describe economic scalability by measuring system efficiency and capacity as a function of required resources. Efficiency represents the amount of resources needed to deliver a unit of service—such as a video stream or e-mail message—or to complete an e-commerce transaction. Capacity represents the maximum rate of service that a system can handle. One architecture has better economic scalability than another if its efficiency and capacity functions grow faster.

As Figure 1 shows, we can use the same metrics to capture the related concept of *system scalability*. A system is considered scalable if capacity continues to grow, even if slowly, as more resources are put into the system. A system is *not* scalable if there is a point at which adding resources yields no increase in capacity.

The difference between system and economic scalability lies in expectations. Computer scientists and engineers have defined ideal system scalability to be constant efficiency and a linear rate of change in capacity, illustrated by the dotted line in Figure 1. An ideally scalable system would grow in capacity in direct proportion to the amount of resources dedicated to it, but no more. Even though certain systems, in practice, exhibit greater scalability, the professional community has generally discounted such cases on the assumption that they did not properly account for all resources.

Vipin Kumar and others proposed one of the best metrics for system scalability, the isoefficiency function (*IEEE Parallel and Distributed Technology*, vol. 1, no. 3, Aug. 1993, pp. 12-21). This metric characterizes the increase in raw workload a system requires to maintain constant efficiency as more resources are added to it. One system is more scalable than another if it has a slower-growing isoefficiency function.

While the isoefficiency framework may be adaptable to the analysis of economic scalability, its name suggests that ideal scalability keeps efficiency constant, which makes economic scalability impossible. A simple computing resources model—processors with fixed computing rates—actually presupposes this impossibility. Thus, we need a richer model to capture the full range of resources contributing to the cost of an Internet-delivered service unit—from cache storage to cable rights-of-way. The model could then support development of an Internet architecture that exhibits economic scalability, not just system scalability.

SUPERCOMPUTING LESSONS

In the late 1980s and early 1990s, the supercomputing industry promised, but failed to deliver, economy of scale. Supercomputers were supposed to solve large problems at a lower unit cost than smaller machines could.

One explanation for the industry's failure focuses on Moore's law, which improved processor and memory component technologies at a rate of 30 to 60 percent per year. A supercomputer that used two-year-old chips suffered a factor-of-two disadvantage in efficiency compared with a smaller system employing the latest chips.

As a result, the supercomputers that survived either use the latest chips in clusters with minimal custom components or run sequential code that offsets system inefficiencies by savings in the human software development effort. Architectures that required special programming and couldn't use the latest chips are now defunct.

The supercomputing industry's experience suggests two strategies to achieve economic scalability for the Internet, one *temporal* and the other *spatial*.

The temporal strategy simply calls for waiting. Moore's law applies in varying degrees to all of the Internet's constituent technologies, except perhaps the costs of cable right-of-ways and satellite launches. By waiting, the Internet becomes cheaper. But we might have to wait a decade or more for the cost of a video stream to drop by the estimated factor of 100 needed to deliver one hour of 1-Mbps video for US$0.05 instead of the current $5.

The spatial strategy concentrates on improving efficiency for a given service as it grows in popularity. This strategy enhances the Internet's commercial potential by promising lower costs for more popular content and applications. By comparison, the temporal strategy improves efficiency for all types of content, applications, and services, whether they are popular or not.

IMPLEMENTATION TECHNOLOGIES

Bandwidth expansion devices and multicast content delivery networks are two technologies that can help.

Bandwidth expanders

Given that bandwidth represents the dominant cost factor in any Internet-based application or service, numerous products evolved over the past decade to expand bandwidth. Caching and compression are key technologies underlying many of these products.

We can analyze a bandwidth expander's effect on a single link. This analysis adopts the viewpoint of an ISP selling bandwidth over an access link, either to a server in the context of a

hosting service or to a business or home.

The exponential growth in total Internet backbone traffic for the past 20 years reflects both more links and more capacity in individual links. For simplicity, assume that these two factors contribute equally to the growth—about 40 percent per year each.

If a bandwidth expander can double a link's capacity, an ISP can postpone upgrading the link for two years. The link's growth rate, however, will remain unaffected. Two years after installing the expander, the ISP will still have to double the link bandwidth and the bandwidth expander's capacity as well.

A skeptic might conclude that eliminating one doubling of a link's bandwidth has negligible effect on Internet efficiency, since the costs saved in the first two years of deploying a bandwidth expander are offset by the costs of running, managing, and upgrading it forever.

However, the technology underlying bandwidth expanders also improves with time according to Moore's law, assuming that the expander's software does not squander the hardware enhancements. Because Moore's law matches the growth in individual link bandwidth by doubling hardware performance every two years at no extra cost, ISPs that incorporate bandwidth expanders can upgrade the device's capacity at a fixed per-link recurring cost. If that recurring cost is cheaper than the cost of supporting links that are twice as large, the network will be that much more efficient than competing networks that lack expanders.

Bandwidth expanders can thus save the difference between the cost of bandwidth and the cost of the devices themselves. They do not, however, fundamentally improve the functional relationship between capacity (bandwidth delivered) and resources (raw bandwidth required).

Multicast-enabled distribution

Multicast technology has a stronger impact on network efficiency, albeit at the severe cost of requiring all receivers to receive the same transmission at the same time, and sometimes even at the same bandwidth. Multicast routing creates a distribution tree of network links in which the root is the source of the information being multicast and the leaves are the receivers. The resources used are routers (tree nodes) and communication links (tree edges).

Bandwidth represents the dominant cost factor in any Internet-based application or service.

For simplicity, assume that the multicast tree is binary, although the conclusions apply in any case. If we double the multicast delivery rate by doubling the number of receivers, we need to add a new level in the tree. The resources everywhere else in the multicast tree remain constant. The number of added leaf links and nodes approximately equals the number of interior links and nodes, thus doubling the tree's size and cost to double its capacity. So, if all resources are equally valuable, the efficiency remains constant after doubling the delivery rate.

In practice, however, links located nearer to the multicast source tend to be physically much longer, which increases their cost relative to leaf links at the network's edge. Deploying long-distance cables requires more digging, the lasers that drive them are expensive, and so is maintenance—especially if they are marine cables or satellite links. Adding resources only at the network's edge doubles a tree's capacity with far less than twice the resources. This makes multicast an economically scalable technology.

Unfortunately, multicast places constraints on the receivers to receive the same content at the same time, which renders it impractical for the vast majority of Internet applications and content. Multicast schemes that relax this requirement reduce their economic scalability in proportion to the nonoverlap of reception by different receivers.

A better solution is to deploy special storage and computation nodes near the network's edge and have them specialize in receiving multicasts of popular content and applications, replicating the material, and then delivering it on demand over the last hop or two in the multicast tree. This basic architecture describes a multicast-enabled content distribution and delivery network.

Multicast-enabled CDNs are economically scalable. Consider the extreme case in which every local area network contains a copy of a given piece of content that has been multicast to a local content-delivery server. In this case, doubling the audience simply doubles the infrastructure's efficiency because no additional resources are required—assuming that bandwidth and server costs on the LAN are negligible.

Any technology capable of a similar resource-utilization pattern can yield similar benefits. Distributed caching, in which *every* router contains a cache, is a case in point.

Multicast-enabled content distribution and delivery networks, as well as distributed caching, provide clearer economies of scale than localized bandwidth-enhancing technologies such as caching and compression on a single link. The horizon is open, however, for exploring technologies that let the Internet deliver its economic potential to the world at large.

Abdelsalam "Solom" Heddaya is chief technology officer and cofounder of InfoLibria. He is also an adjunct associate professor of computer science at Boston University. Contact him at a@heddaya.net.

Editor: Sumi Helal, Computer and Information Science and Engineering Dept., University of Florida, P.O. Box 116125, Gainesville, FL, 32611-6120, helal@cise.ufl.edu

Whither Warhol's Law?

Wayne Wolf, Princeton University

About 10 years ago, someone asked me to provide one of those brief "expert opinion" boxes for *IEEE Spectrum*. I responded with what I called "Warhol's Law of Computer Systems Architecture." This law stated that every computer architecture would be the price-performance leader for 15 minutes, a reference to artist Andy Warhol's famous comment that, in the future, everyone would be famous for 15 minutes.

I made this comment in jest, but certainly many new CPU architectures—Sparc, MIPS, Alpha, PA-RISC, and so on—were coming on the market at the time and jockeying for position as the leader. The churn was amazing.

I was thinking mainly of workstation CPUs, but I believed the embedded CPU market would follow suit. Embedded systems provide an ideal forum for CPU competition. Car shoppers don't care what computer controls the engine; they just want the car to run well. A lot of people seemed to want to design CPUs, so I reasoned that no end of new competitors would appear to introduce new features and drive down prices.

Every computer architecture will be the price-performance leader for 15 minutes.

THE FUTURE THAT WASN'T

It hasn't worked out that way, however. The high-end market has several players, but not the profusion I expected a few years go. ARM, widely used in the cell phone market, enjoys a huge customer base—given that more than 500 million cell phones were sold last year. Further, Intel bases its StrongARM and XScale processors on the ARM architecture.

MIPS is another architecture that chip designers use for intellectual property. Other architectures that chip manufacturers traditionally provided as IP, such as Motorola's 68000 family, are also strong in the embedded field. But all these architectures have existed for 15 years or more, and all were originally designed with the general-purpose computer market in mind: ARM for the Acorn computer, MIPS for a variety of workstations, and the 68000 as a contender for the IBM PC.

There have been some new entries, particularly on the very-long-instruction-word front. Texas Instruments has launched the C6x VLIW machine, which has become successful as a signal processing engine. The StarCore VLIW machine, offered as IP, is also used in signal processing and networking. The TriMedia VLIW, spun out of Philips and optimized for video processing, has found use in TiVo's personal video recorder.

However, I thought more CPUs would be slugging it out in the embedded space by now. Clearly I was wrong. So why hasn't the embedded marketplace offered more fertile ground for CPUs?

NETWORK EFFECTS

Networked appliances are emerging as one reason for a few CPUs to command market share. Digital music devices are the best example of this phenomenon today. Various devices let users download music from the network or play CDs created on a computer. The music industry, anxious to protect its copyrights, is pushing for digital rights management software. DRM software runs on the playback device to ensure that users pay for the right to play digital copies.

Microsoft is a major supplier of DRM software as part of its Windows Media Player. It supplies only object code, however, not source code, and severely restricts the architectures on which it offers code support. This is understandable, given that Microsoft must try to protect its security protocol while making sure the code works correctly. But it also tends to push device designers to those platforms that support the required software.

We can expect the same pattern in networked video devices and any others that connect to the Internet. Java mitigates the ties between platforms and applications to some extent, but embedded systems use Java primarily in non-real-time functions. Applications tied to real-time or low-power performance show up on only a few platforms.

TOOLS AND I/O

More traditional reasons for the few consistent suppliers in the embedded CPU market run parallel to the reasons

for architectural convergence in the desktop and server spaces.

First, a CPU is only as useful as the toolset backing it up: compilers, assemblers, linkers, loaders, and the rest of the programming chain. The toolset also includes debuggers and even libraries for important functions, such as digital filtering and I/O device support.

It takes years to develop and debug these tools. Developers generally don't make that effort unless the target CPU commands a substantial market. One reason for the 68000 family's popularity in embedded computing—Ford was a bigger customer than Apple for 68000s—was its strong toolset from the Macintosh.

Similarly, the range of devices that attach to the CPU factors into the choice. Specialized devices and interfaces for a particular project may require custom design, but many devices are common across projects: timers, direct-memory-access controllers, general-purpose I/O, and so on. Few embedded system designers have either the time or the inclination to recreate them from scratch. As in the desktop market, a CPU that doesn't support existing devices would never get off the ground. Likewise, successful CPUs tend to have more I/O devices designed for them, which attracts more customers and yet more opportunities to add to the line of available devices.

COSTS OF CHANGE

Integrated-circuit manufacturing can pose a significant entry barrier for a system-on-chip. Generally, SoCs must sell a million units to pay off the mask costs and other design costs associated with deep submicron IC technology. Many designers are reluctant to be first on a given fabrication line. They would prefer to know that the hidden problems—timing bugs, power distribution problems, and so on—have been worked through already. As a result, CPU IP vendors work with the major chip foundries to qualify CPU designs, fabricating and testing a design to make sure it will work properly. If the vendor supplies a family of CPUs, then it must qualify all members; if the foundry goes to a new fabrication process, the vendor must requalify them on the new fab line.

This is a lot of work, not to mention expensive. Foundries seldom want to qualify a CPU unless they believe the market for it already exists. As a result, successful CPU architectures tend to stay successful.

Cultural reasons contribute to the ongoing popularity of embedded CPUs as well. Designers tend to use components they are familiar with, and a company that has invested in the tools for a particular processor has a strong reason to continue using it in subsequent designs.

Configurable processors allow a head-spinning number of architectures.

Nor do many designs start with a clean sheet of paper. Legacy hardware and software often make it much more attractive to base new designs on existing CPUs instead of retooling for another processor.

CONFIGURABLE PROCESSORS

One new trend, however, is putting an interesting spin on the embedded CPU. Several companies now offer processors with configurable architectures. Because the architecture itself—not its logic implementation—is reconfigurable, these processors don't have to be built on a reconfigurable device like a field-programmable gate array. Designers can adapt the CPU to the application at hand—for example, adding specialized instructions to speed up critical operations and new function units in the data path to give these instructions extra oomph.

Configurable processors are challenging. They require the same level of support as any other CPU—compilers, assemblers, linkers, debuggers—but the tools must be configurable along with the CPU. And that configurability must occur consistently across the CPU and all the tools. It's hard enough to get a fixed CPU design right. Imagine getting it right when someone else can add to the instruction set.

Configurable processors allow a head-spinning number of architectures. Not only can every customer design a different instruction set, but a single chip could contain several CPUs, each configured for a different instruction set. In a complex application that requires several process stages, using multiple CPUs, each with different instruction sets, provides not only task-level parallelism but also instruction-level speedups for each task.

Vendors must necessarily synthesize configurable processors from a hardware description language model. They can use standard analysis tools to ensure that the CPU hardware runs at the required speed. It would be impossible to qualify every variation of a configurable processor, but designers seem comfortable using these processors in chips—probably because the manufacturer bases all the CPU variants on a common description. Similarly, because developers customize the compilers and other tools from a known source instead of building them from scratch, designers can leverage the tool development over many CPU instantiations.

Configurable CPUs let all sorts of people, not just computer architects, experiment with new instruction sets. I'm sure that designers are experimenting with new instructions right now. Some of them won't pan out, but some will. Perhaps every instruction set will get its 15 minutes of fame after all. ■

Wayne Wolf is a professor of electrical engineering at Princeton University and author of Computers as Components: Principles of Embedded Computing System Design *(Morgan Kaufman, San Francisco, 2000). Contact him at wolf@princeton.edu.*

Web Services and Context Horizons

Clay Shirky

Web services attempt to make interapplication communication as easy as the Web made publishing documents. By defining a set of structured, human-readable data formats and making them open and independent of any particular language or platform, Web services designers hope to create a framework for program-to-program interoperability. This should allow binding Internet-accessible services together in various combinations even though different parts of the application are distributed over the network.

Loose coupling vastly increases the scale of deployable systems. In its most ambitious version, the Web services framework functions like an Internet-scale operating system, assembling complex programs ad hoc from pieces anywhere in the world. This system would use the Internet as a kind of global bus akin to the local bus that ties together an individual computer's components.

General interoperability of software across networks is a goal with a long history but only partial successes, such as the Distributed Component Object Model (DCOM), common object request broker architecture (Corba), and remote method invocation (RMI) over the Internet inter-ORB protocol (IIOP). The question is whether Web services' contribution to interoperability will be revolutionary or only incremental.

For Web services to succeed, software engineers must rethink what the terms "local" and "global" mean in an Internet-scale operating system.

THE CHALLENGE

Web services cannot yet support an Internet operating system. In their current state, they are simply plumbing for the exchange of XML documents using SOAP (the simple object access protocol)—http://www.w3.org/TR/SOAP/.

The most widely known example of a Web service is the Google API (http://www.google.com/apis/), which lets remote applications send search requests to the Google search engine packaged as a SOAP call. Google executes the request and returns a structured XML document containing the results to the calling program. This lets software designers build Internet searching directly into an application, without having to write a custom screen-scraper function to extract Google results from an ordinary Web page.

More interesting possibilities lie in using Web services to integrate multiple functions. For example, consider a London-based firm that wants to automate purchasing. Upon execution of an initial request for the purchase of 1,000 units of some item, the resulting query passes over the network to a Web service run by the supplier in Germany. That supplier in turn calls a Web service that calculates shipping cost, a second that computes tax, and a third that converts between pounds and euros. The German supplier then returns a consolidated quote to the original caller in London.

An application assembled in this way can consist of many pieces—some running on the same computer as the calling program, some running on a local network, and others running on the Internet—and any of those pieces can contain still more remote pieces. A single request can trigger a cascade of subrequests until it becomes impossible to design a Web services application that knows in advance all the parts it will ultimately use.

THE NETWORK IS NOT THE COMPUTER

Web services would seem to usher in the world long ago predicted by Sun Microsystems, where "the network is the computer." However, as catchy as that phrase is, it doesn't accurately describe networked services because local-scale solutions for problems such as latency, reliability, and security don't work at Internet scale.

Traditionally, the distinction between local and global resources has been relatively simple. On a machine, local is whatever is in the same chassis—CPU, storage, memory—and an application running locally usually has implicit access to those resources. On a network, machines that share a local context are typically inside a shared boundary, and an array of firewalls, proxy servers, and other special-pur-

pose computers enforce that boundary. Thus, local is whatever is "in here"—whether a particular machine, LAN, or intranet—and global is everything else.

The Web services framework disregards this separation. Rather than a computer that runs several applications, Web services offer the possibility of running an application on several computers, many of which the user neither owns nor controls. Every Web services interface can offer not only its own services, but also those it knows about, even if they reside elsewhere.

Web services require defining a new meaning for local, in which different kinds of applications operate within different network horizons. Understanding these *context horizons*, and creating the tools and techniques for managing them, are essential to achieving an Internet-scale operating system.

CONTEXT HORIZONS

There are many kinds of context horizons, some specific to individual applications. Financial institutions, for example, require various atomic transactions that are difficult to link together via the kind of "loose coupling" that the Web relies on. Several groups, including the Organization for the Advancement of Structured Information Standards (OASIS; http://www.oasis-open.org/who/), are designing protocols for this kind of transactional support.

Some context horizons are of such general importance, though, that they will profoundly affect the overall deployment of Web services and whether these services actually become part of an Internet operating system or simply continue plumbing for XML. Three of the most important are trust horizons, semantic horizons, and coordination horizons.

Trust horizons

If an application calls a remote Web service that in turn calls yet another Web service, how can you be sure that the calls do not compromise security or privacy requirements?

Trust has two components. *Outbound* trust hinges on private information not passing into untrustworthy hands—if A calls B and B calls C, A must be able to trust C as well as B, even if A doesn't know C exists. *Inbound* trust depends on remote sources of information providing trustworthy results—if A gets data from B, who gets it from C, A must trust C as well, even though C is invisible to A.

A given Web service might not even know how many different applications it is part of.

Thus, unlike the one-to-one matching between individual users and machines that typically occurs today, trust horizons will require more subtle constructions ("I trust this Web service with my Web site's address but not with my Social Security number") and a more complex handling of transitive trust ("I trust this Web service with my Social Security number, but I don't trust it to forward that number elsewhere" or "Any additional Web services involved in processing my data must sign their replies so I can generate an audit trail").

The obvious tools for handling trust horizons—a set of identities for people, machines, and transactions capable of validation and tracking—do not yet exist, and the concept is surrounded by contentious privacy and control issues. As Web services grow, the imperative to manage increasingly complex trust horizons will drive development of Internet-scale identity and authorization systems.

Semantic horizons

Web services' current state is analogous to international snail mail. The ability to send a letter from one country to another does not guarantee that the recipient will be able to read it. Likewise, Web services create a method for exchanging data and requests between remote applications, but they cannot ensure that the remote application receives understandable data.

The problem is lack of a global ontology, a single framework for describing everything. A seemingly simple phrase such as "catalytic" means one thing to an automotive engineer, another to a biochemist, and yet another to a business consultant. Because Web services must support numerous user groups, a request passing from one context to another with a subtly altered meaning will be a constant possibility.

Indeed, the ease of creating XML documents will motivate more people to propose standards. Sorting out the semantic collisions, and determining how to write applications that know what to do when they are talking to remote Web services that do not share their semantic scope, represent significant challenges.

Coordination horizons

There is no limit on the number or type of Web services that could be woven together to create a particular application, nor is there a limit to the number of applications a given Web service can participate in. Depending on the handling of trust horizons, a given Web service might not even know how many different applications it is part of, and a given calling program might not know how many Web services are participating in a given request.

This introduces a dilemma that has plagued software designers for decades, namely, how large should an application be? The Unix philosophy suggests that applications should be small and have generic interfaces, while the personal computer philosophy is that applications should contain all the functions they need internally and have customized interfaces. Web services seem to favor the Unix philosophy.

However, the Web services framework encourages customized rather than standardized interfaces. Further, many of the standard techniques for designing componentized software, such as inheritance and overloading,

won't work with Web services because the remote service is opaque and can be accessed remotely, but not copied or modified locally.

Thus, the question of whether to add a particular feature or function to an existing Web service or create a new service and interface is a critical design issue, and solutions developed over years of modular design of local software no longer apply. It will be critically important for Web services to operate within their particular coordination horizon, which will itself be affected by modularization, network issues such as latency and reliability of remote routes and nodes, and other context horizons such as trust and semantics.

SIZE MATTERS

In the 1930s, economist Ronald Coase questioned why firms are the size they are. If the market can provision everything a firm needs, he wondered, why do businesses not simply outsource every function? The answer, Coase concluded, was transaction costs.

Theoretically, it's possible to create a virtual company by contractually handling every function, from HR to printing to R&D, but designing and executing contracts carry a cost. Coase's intuition was that the firm would grow to the size it needed to be to minimize contracting costs for essential services by bringing those functions into one organizational context.

This is analogous to the main problem facing Web services. With a SOAP interface, every function or method call could pass over a network, but the program would slow to a crawl. Yet if every program monolithically encapsulates everything it needs, the Web services framework offers no advantage. Web services must be large enough to exploit remote resources without becoming so dissipated that they cease working.

Given the context horizons implicit in Web services, transaction costs will profoundly impact application design. Choices about the amount and type of encapsulation within a given Web service will in turn affect overall issues of size and coordination among interoperating services. To succeed even moderately, Web services will require a new software engineering philosophy that redefines what the terms local and global mean in a large-scale networked environment.

Clay Shirky is a writer and consultant on Internet technologies. Contact him at clay@shirky.com; http://www.shirky.com.

Editor: Richard G. Mathieu, Dept. of Decision Sciences and MIS, St. Louis University, St. Louis, MO 63108; mathieur@slu.edu

Continued from page 104

sibility for the work of those skilled in particular trades:

- civil engineers oversee machine drivers and concrete workers,
- mechanical engineers oversee machinists and other metalworkers, and
- electrical engineers oversee riggers and electricians.

In traditional engineering, professional work is clearly and formally distinct from construction work and from the work the constructed products' users do. Not so in the computing profession. Where segregation of this kind exists, it is local and informal. The computing profession seems instead to want to separate into branches, each of which engineers a specific type of system: information, computer, software, knowledge, and so on.

Industry generally sees the computing profession as a variegated and vacillating collection of leaderless skills. That's why industrial leaders can rant about IT skills with little fear of contradiction—they see no clear structure because the profession itself sees none.

Engineers and tradespeople have distinct and essential roles. Engineers lead by bringing their education and experience to bear when applying general principles to new problems. Tradespeople carry out assigned tasks using a variety of tools and techniques coupled with skills acquired through training and experience.

General principles are stable and only slowly extended, and engineers apply those principles continually to develop new tools and techniques for using them. Thus, training and continual retraining are much more important for the tradespeople who use those new tools and techniques than for the engineers who develop them. In many countries, specialist schools for technical or vocational education supply this training. If they were better supported in specialist areas of computing by the profession, and if computing professionals were prepared to lead computing tradespeople, any real shortage of particular IT skills would be short-lived.

If the computing profession recognized and encouraged the development of computing trades, we would have little need for separating the profession into branches, and the education of computing professionals could be redeveloped along generalist lines, perhaps even as a unified data engineering discipline. Let's face it, programming is a craft and trade, not a profession.

Although we once could imagine a profession distinguished from others by its exclusive use of digital computers, this notion is now ludicrous.

The secondary profession

A data engineering generalist would seek to exploit data. As a resource, data and the machines that process data differ in kind from the resources and machines that other engineering branches exploit. Engineering's traditional branches exploit resources and phenomena derived from them, such as structural materials and kinetic and electrical energy.

Data, on the other hand, subsists in immaterial representations *imposed on* material derived from natural resources. Data is thus a secondary, indirect, and limitless resource that all professions and occupations use. Indeed, humanity has based its civilization on the formal use of data.

Although we once could imagine a profession distinguished from others by its exclusive use of digital computers, this notion is now ludicrous. Yet computing professionals continue to talk loosely of the *computer profession*, a phrase echoed in the "Innovative Technology for Computer Professionals" tagline that, unfortunately, appears on this publication's front cover.

All professions—and those who educate their practitioners—will increasingly depend on machines based on digital technology. Therefore, educators should teach future computing professionals more than data engineering. They should educate students to accept and promote their profession as secondary: one that aids and abets other professions. The students' education should include project work carried out in cooperation with students from other professions. To survive and flourish, the computing profession must abandon the idea of living in splendid and oracular isolation.

We *do* need a branch of engineering devoted solely to the design and manufacture of digital computers, which might well be called computer engineering. But computer manufacturing's oligarchic nature means that the field will need relatively few such computer engineers, so their professional education and affiliation would be better regarded as a specialization within electronic engineering.

BEYOND PROFESSIONS

Digital machinery is rapidly coming to dominate upper- and middle-class domestic life in developed countries, most significantly in entertainment applications. In such an environment, children receive a significant amount of computer and Web training during their daily activities. Many occupations will rely increasingly on using computers and computer-based machinery, and training in the vocational use of computers and such machines—training distinct from that of computing tradespeople—should be demanded by the relevant trade organizations.

Using appliances based on digital computers lies beyond the computing profession's proper concerns. Using computers to compute does not. The profession has a responsibility to promote and support the effective and knowledgeable vocational and domestic use of computation. It should thus give attention to how marketing objectives dominate such computation.

Ordinary users and their teachers feel flummoxed by the ever-growing complexity of the personal computers, oper-

ating systems, and software suites they must upgrade continually. If the automotive industry behaved as the computer industry does, it would sell in place of the private car an articulated 18-wheeler containing a bathroom and private theater, with an automatic transmission supplemented by a 233-speed manual gearbox necessary for curves and hills. We desperately need a program suite that combines *basic and stable* spreadsheet, database, document, and graphic processing under a *basic and stable* operating system on a *basic and stable* personal computer. With such a foundation, *all* children and apprentices might effectively learn persistently useful skills. The digital divide might then start narrowing rather than continuing to widen. Further, if we could make available *basic* RPG, Cobol, and Fortran in the simple style of the 1960s, many more young people might learn to do their own programming.

All this talk of IT skills masks a deception: The phrase itself suggests that such skills form some kind of esoteric capability distinct from all others. The computing profession should assert most vigorously that everyone should possess the skills to deal with data effectively. Basic education should have this aim. After all, we expect schools to inculcate literacy and numeracy in their students—skills essential to IT. In this age and in the developed world, the essential aspects of literacy and numeracy should include acquiring basic competency in document and spreadsheet manipulation.

Vocational education should also aim to impart skills in the literate and numerate use of computation within every trade. Schools should avoid vocational training in IT skills per se. Rather, the aim should be to produce workers in specialist fields, such as computer security and programming, who can work with professionals in various fields, but particularly with computing professionals.

If we accept computing as a secondary profession, and if we view all professions and vocations as benefiting from skillful data use, people seeking *professional* jobs will see that having computers and being able to program them matters "jack squat" compared to being able to help other people use them and benefit from that use. With luck, employers and government will come to see this, too. ■

***Neville Holmes** is an honorary research associate at the University of Tasmania's School of Computing. Contact him at neville.holmes@utas.edu.au. Details of citations in this essay and links to further material are at http://www.comp.utas.edu.au/users/nholmes/prfsn.*

Jobs, Trades, Skills, and the Profession

Neville Holmes, University of Tasmania

Recently, I got an e-mail message from someone who assumed I wrote the June column ("Using Computers in Our Daily Life," *Computer*, June 2002, pp. 104, 103). This person took issue with the statement that those who "remain tied to the past and refuse to adopt digital technology ... can expect nothing but poverty." That Ana Asuaga, the June column's author, lives in Uruguay, a Third World country, gave a particular edge to that assertion.

This e-mail came from the US, where its author benefits from a First World environment and skills: He writes his own software, maintains Web pages, exposes security holes on the Internet, and rescues computers from dumpsters so that he can refurbish them. Yet he cannot get a job in the computing field, complaining bitterly that "just because you have a ton of computers and can program them doesn't mean jack squat."

IT SKILLS SHORTAGE

Such a plaint might, in the absence of further information, suggest that some aspect other than talent and experience makes this person unemployable. However, shortly after receiving this message, I chanced on a column by the highly respected Australian technology journalist Graeme Philipson ("IT Skills: A Shortage or a Scam?" *The Age*, 18 June 2002; http://www.theage.com.au/articles/2002/06/15/1023864366686.html). Philipson sees the IT skills shortage as nothing less than a political scam.

Is the reputed shortage of IT skills a political scam or is the computing profession itself at fault?

So do many of his readers, who contributed to a response he describes as by far the greatest he has experienced in 20 years of writing about IT ("Testimonials Put the Lie to the Myth of the IT Skills Shortage," *The Age*, 16 July 2002; http://www.theage.com.au/articles/2002/07/15/1026185154255.html). All but one of his more than 250 respondents agreed with Philipson—those he quoted echoed the sentiments of my own irate reader.

Nor is this solely an Australian phenomenon. Certainly, Norman Matloff's testimony to the US House Judiciary Committee's Subcommittee on Immigration ("Debunking the Myth of a Desperate Software Labor Shortage"; http://heather.cs.ucdavis.edu/itaa.real.html) indicates that the situation is much the same in the US. Further, the main government measure to counteract the supposed shortage—encouraging people with IT skills to emigrate from less-developed countries—is common to both Australia and the US, and it also seems to be popular in Western Europe.

Who is behind this myth? In his June column, Philipson suggested that "many in the IT industry are manufacturing fears of an IT shortage to get handouts from government and to be able to hire cheaper immigrant labour." And in the US, an H-1B Hall of Shame lays blame similarly (http://www.zazona.com/ShameH1B/).

Blaming the IT industry or a servile government for this deceit is unfair. The main fault lies with the computing profession.

IT AND THE PROFESSION

Information technology concerns itself with processing data, particularly digital data. The computing profession takes this activity to be its special bailiwick.

But what kind of a profession are we? We can lay some claim to being a branch of engineering. The various engineering branches exist to exploit technically physical materials and other resources. Civil engineering exploits materials in static structures, mechanical engineering exploits kinetic energy, and electrical engineering exploits electrical energy. The computing profession exploits data: conventional representations of facts or ideas.

The computing profession and traditional engineering branches differ in two respects: their professionals' responsibilities and the nature of the resources they exploit.

Professional responsibility

In traditional branches of engineering, professionals lead and take respon-

Continued on page 102

WWW2003

2003 WWW Conference

Programme Committee Co-Chairs

Yih-Farn Robin Chen, AT&T Labs-Research
László Kovács, MTA SZTAKI
Steve Lawrence, NEC Research Institute

Tutorials/Workshops Chair

Michael Bieber, New Jersey Institute of Technology
Beatrix Toth, MTA SZTAKI

Developers Day Co-Chairs

Ivan Herman, W3C
János Szël, MÁV Informatics

Conference Co-Chairs

Gusztáv Hencsey, MTA SZTAKI
Bebo White, Stanford Linear Accelerator Center

IW3C2 Liason to WWW2003

Ivan Herman, W3C

WWW2003 Organisers

International World Wide Web Conference Committee (IW3C2)
http://www.iw3c2.org

Computer and Automation Research Institute of the Hungarian Academy of Sciences (MTA SZTAKI)
http://www.sztaki.hu

WWW2003 Conference Partner

International Federation for Information Processing (IFIP)
Working Group 6.4 on Internet Applications Engineering
World Wide Web Consortium

The International World Wide Web Conference Committee (IW3C2) and the Computer and Automation Research Institute of the Hungarian Academy of Science (MTA SZTAKI) cordially invite you to participate in the 12th World Wide Web Conference in Budapest, Hungary, 20-24 May 2003.

Beginning with the first International WWW Conference in 1994 this prestigious series of the International World Wide Web Conference Committee (IW3C2) also provides a public forum for the WWW Consortium (W3C) through the annual W3C track.

The conference will consist of a three-day technical program, preceded by a day of tutorials and workshops and followed by a Developers Day. The technical program will include refereed paper presentations, alternate track presentations, plenary sessions, panels, and poster sessions describing current work. The tutorials and workshops will provide in-depth looks at specific areas of current interest. Developers Day will be devoted to in-depth technical sessions designed specifically for Web developers.

The Twelfth International World Wide Web Conference

Budapest Convention Centre 20-24 May 2003 Budapest, HUNGARY

REFEREED PAPERS

- Applications
- Browsers and User Interfaces
- E-Commerce
- Hypermedia
- Mobility and Wireless Access
- Multimedia
- Performance and Reliability
- Search and Data Mining
- Security & Privacy
- Semantic Web
- Web Engineering

ALTERNATE TRACKS

- Education
- Global Community
- Industrial Track
- Practice & Experience
- Web Services
- W3C Track
- Panels

Key Deadlines

Paper Submission :	November 15, 2002
Author Notification :	January 31, 2003
Final Paper:	February 28, 2003

http://www2003.org

If you have any questions or wish to be placed on the WWW2003 announcements mailing list, please send email to info@www2003.org

IFIP W3C MTA SZTAKI IW3C2

EXPOSE.DEBUG.PUBLISH.

Expose errors at runtime with Parasoft Insure++®.
Debug applications quickly. Publish with greater efficiency and profitability.

Parasoft Insure++ is an automated runtime error detection tool that keeps hard-to-find errors from creeping into your finished product. From memory corruption, memory leaks and memory allocation errors to large classes of coding and C++™ specific errors, Insure++ lets you find elusive defects and fix them fast.

Effortless runtime testing. Accelerated development cycles. Increased profit margins.

With Insure++ runtime error detection is efficient and foolproof. No time wasted. No errors undetected. So your applications hit the market faster, and don't bounce back once they get there.

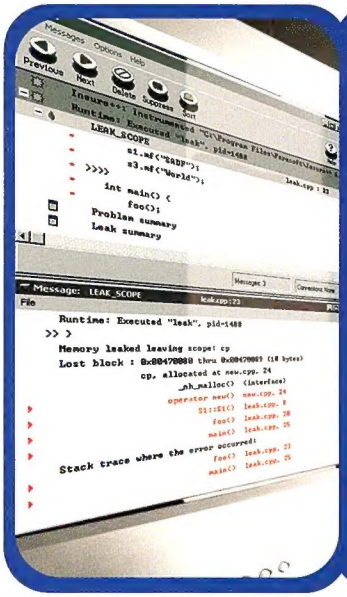

Features

Finds errors other tools overlook.

Provides information necessary to fix errors.

Builds a database of program elements, and checks values against it at runtime.

Checks third-party libraries and functions.

Checks all types of memory references.

Benefits

Streamlines error detection.

Reduces development and support costs.

Increases understanding of memory usage.

Requires no modification of source code.

Provides flexible levels of error detection.

Platforms

Windows NT/2000/XP

Linux

Tru64Unix (4.x, 5.x)

IBM RS/6000 (AIX 4.3.x, AIX 5L)

HP (HP_UX 10 & 11)

SGI (IRIX 6.5)

Solaris SPARC (6, 7, & 8)

Solaris x86 (7 & 8)

AEP™
AUTOMATED ERROR PREVENTION

Parasoft Insure++
Automated Error Prevention for C/C++™

To download an evaluation copy of Insure++ go to http://www.parasoft.com/com9.
Or call a Parasoft Error Prevention Specialist at (888) 305-0041.

PARASOFT insure++®